ArubaOS 8.7.0.0



a Hewlett Packard
Enterprise company

User Guide

Copyright Information

© Copyright 2024 Hewlett Packard Enterprise Development LP.

Open Source Code

This product includes code licensed under the GNU General Public License, the GNU Lesser General Public License, and/or certain other open source licenses. A complete machine-readable copy of the source code corresponding to such code is available upon request. This offer is valid to anyone in receipt of this information and shall expire three years following the date of the final distribution of this product version by Hewlett Packard Enterprise Company. To obtain such source code, send a check or money order in the amount of US $10.00 to:

Hewlett Packard Enterprise Company
6280 America Center Drive
San Jose, CA 95002
USA

Contents 3

Revision History 15

About this Guide 18
- New Features and Hardware Platforms 18
- Fundamentals 24
- JSON APIs 27
- System Requirements 27
- Supported Browsers 27
- Related Documents 28
- Conventions 28
- Contacting Support 29

Mobility Master Configuration Hierarchy 30
- Understanding Configuration Hierarchy 30
- Centralized Configuration 32
- Configuration Validation 37
- Serviceability 37
- Bad Configuration Recovery 38
- Disaster Recovery 39
- Initial Provisioning Recovery 39
- Mobility Master User Interface 39

MultiVersion Support 44
- Important Points to Note 44
- UI Support for Multiversion 44

The Basic User-Centric Networks 46
- Understanding Basic Deployment and Configuration Tasks 46
- Managed Devices Configuration Workflow 49
- 7200 Series Controllers Port Behavior 51
- Using the LCD Screen 51
- Configuring a VLAN to Connect to the Network 54
- Configuring User-Centric Network 61
- Replacing a Controller 61

Control Plane Security 66
- Control Plane Security Overview 66
- Configuring Control Plane Security 67
- Managing AP Whitelists 68
- Whitelist DB Optimization 76
- Configuring Networks with a Backup Mobility Master 76
- Replacing a Controller on a Multi-Controller Network 77
- Troubleshooting Control Plane Security 78

Network Configuration Parameters 80

Campus WLAN Workflow 80
Understanding VLAN Assignments 81
Configuring VLANs 83
Trusted and Untrusted Ports and VLANs 91
Assign an IP Address to a VLAN 92
Configuring Trusted or Untrusted Ports and VLANs 96
Configuring the Mobility Master IP Address 98
Configuring the Loopback IP Address 99
Configuring Static IP Routes 99
GRE Tunnels 100
GRE Tunnel Groups 107
Jumbo Frame Support 108
PVST+ 110
RSTP 111
PortFast and BPDU Guard for Spanning Tree 113
LLDP 115

Port Channel Link Aggregation Control Protocol 119
Configuring Port Channel LACP 119
LACP Best Practices and Exceptions 120
LACP Sample Configuration 121

IPv6 Support 122
Native IPv6 Support 122
Supported Applications 123
Important Point to Remember 124
Enabling IPv6 125
Enabling IPv6 Support for Mobility Master and APs 125
Filtering an IPv6 Extension Header 134
Configuring a Captive Portal over IPv6 135
Working with IPv6 RAs 135
Centralized Licensing Support for IPv6 140
IPsec Support 143
RADIUS Over IPv6 153
TACACS Over IPv6 154
DHCPv6 Server 154
Understanding ArubaOS Supported Network Configuration for IPv6 Clients 160
Understanding Authentication and Firewall Features that Support IPv6 161
Understanding IPv6 Exceptions and Best Practices 166

OSPFv2 168
Important Points to Remember 168
Understanding OSPFv2 by Example using a WLAN Scenario 169
Understanding OSPFv2 by Example using a Branch Scenario 170
Configuring OSPF 171
Exporting VPN Client Addresses to OSPF 171
Sample Topology and Configuration 172

Tunneled Nodes 184
Understanding Tunneled Node Configuration 184
Configuration Procedures 185
Dynamic Segmentation 186

Authentication Servers 192
Configuring Authentication Servers and Server Groups 192
Understanding Authentication Server Best Practices and Exceptions 192
Understanding Servers and Server Groups 192

Configuring Authentication Servers ... 193
Configuring an LDAP Server ... 204
Configuring a TACACS+ Server ... 205
Configuring a Windows Server ... 206
Managing the Internal Database ... 207
Configuring Server Groups ... 208
Assigning Server Groups ... 214
Configuring Authentication Timers ... 219
Authentication Server Load Balancing ... 220
Testing a Configured Authentication Server ... 220

MAC-Based Authentication ... 222
Configuring MAC-Based Authentication ... 222
Configuring Clients ... 223
Multi Pre-Shared Key ... 224

Managed Devices at Branch Offices ... 226
Learn more about Managed Device Optimization ... 226
Managed Device Feature Overview ... 226
Zero-Touch Provisioning Overview ... 228
WAN Authentication Survivability Overview ... 232
Managed Device WAN Dashboard ... 236
Using ZTP to Provision a Managed Device ... 237
Using ZTP with DHCP to Provision a Managed Device ... 242
Health Check Services for Managed Devices ... 244
WAN Optimization Through IP Payload Compression ... 245
WAN Interface Bandwidth Priorities ... 246
Uplink Monitoring and Load Balancing ... 247
Hub and Spoke VPN Configuration ... 252
IP Routes Configuration ... 254
Uplink Routing using Next-hop Lists ... 255
Policy Based Routing ... 257
Address Pool Management ... 259
Configuring WAN Authentication Survivability ... 263
Preventing WAN Link Failure on Virtual APs ... 264
Managed Device Integration with a Palo Alto Networks Portal ... 265

802.1X Authentication ... 268
Understanding 802.1X Authentication ... 268
Configuring 802.1X Authentication ... 271
Configuring and Using Certificates with AAA FastConnect ... 277
Configuring User and Machine Authentication ... 278
Working with Role Assignment with Machine Authentication Enabled ... 278
Enabling 802.1X Supplicant Support on an AP ... 280
Example Configurations ... 281
Performing Advanced Configuration Options for 802.1X ... 299
Application Single Sign-On Using L2 Authentication ... 300
Device Name as User Name for Non-802.1X Authentication ... 302
Enhanced Open Security ... 303
Support for WPA3 ... 306

Stateful and WISPr Authentication ... 311
Stateful Authentication ... 311
WISPr Authentication ... 312
Stateful Authentication Best Practices ... 312
Configuring Stateful 802.1X Authentication ... 312

Configuring Stateful NT LAN Manager Authentication 313
Configuring Stateful Kerberos Authentication 315
Configuring WISPr Authentication 316

Certificate Revocation 319
Understanding OCSP and CRL 319
Configuring the Mobility Master or Managed Device as an OCSP Client 320
Configuring the Mobility Master or Managed Device as a CRL Client 321
Configuring the Mobility Master or Managed Device as an OCSP Responder 323
Certificate Revocation Checking for SSH Pubkey Authentication 324

Captive Portal Authentication 327
Mobility Master-Managed Device 327
Stand-alone Controller 327
Understanding Captive Portal 327
Configuring Captive Portal in the Base Operating System 329
Configuring Captive Portal with a PEFNG License 331
Sample Authentication with Captive Portal 334
Configuring Guest VLANs 343
Configuring Captive Portal Authentication Profiles 344
Enabling Optional Captive Portal Configuration 349
Personalizing the Captive Portal Page 353
Creating and Installing an Internal Captive Portal 357
Creating Walled Garden Access 366
Enabling Captive Portal Enhancements 368

Controller Clustering 374
Requirements 374
Key Consideration 374
Support for Homogeneous Cluster 375
Support for Heterogeneous Cluster 375
Cluster Connection Types 376
Roles 377
Remote AP Support 378
IPv6 Cluster Support 378
Cluster Features 379
Authorization Server Interaction 381
AP Failover to Different Cluster 383
Grouping Managed Devices Within a Cluster 383
AP Node List 383
APmove 383
EST Support for Cluster 385
Remote AP Support with Cluster behind NAT 385
VRRP ID and Passphrase 386
Cluster Configuration 387
Cluster Load Balancing 392
Cluster Deployment Scenarios 397
Upgrading Cluster 401
Troubleshooting Cluster 405

Dashboard Monitoring 410
Dashboard Pages 410
Overview 411
Infrastructure 421
Traffic Analysis 430
Security 442

Services 448

MultiZone 453
Functional Flow of a MultiZone AP 453
Important Points 454
Licenses for MultiZone 454
Hybrid CPsec, Mesh AP, and Mobility Controller Virtual Appliance Support 454
AP LACP Support for MultiZone 455
Client Match Support for MultiZone 455
RSDB and Dual 5G Bands Support for MultiZone 455
Configuring MultiZone 456

Virtual Private Networks 459
Planning a VPN Configuration 459
Working with VPN Authentication Profiles 463
Configuring a Basic VPN for L2TP/IPsec 465
Configuring a VPN for L2TP/IPsec with IKEv2 471
Configuring a VPN for Smart Card Clients 476
Configuring a VPN for Clients with User Passwords 477
Configuring Remote Access VPNs for XAuth 479
Working with Remote Access VPNs for PPTP 480
Working with Site-to-Site VPNs 480
Working with VPN Dialers 490

Roles and Policies 491
Firewall Policies 491
Creating a User Role 501
Workflow for Assigning a User Role 504
Understanding Global Firewall Parameters 509
AppRF 2.0 515
Workflow for configuring AppRF 2.0 516
Using an Exclude List 520
Debugging 523
Netdestination and Netservice Aliases 524
IP Classification-based Firewall 526

ClearPass Policy Manager Integration 529
Introduction 529
Important Points to Remember 529
Enabling Downloadable Role on a Managed Device 530
Sample Configuration 531
Per-Command Authorization for Management Users 535

Configuring WLANs 536
Basic WLAN Configuration 536
WLAN Configuration Profiles 540
Configuring the Virtual AP Profile 543
Radio Resource (802.11k) and BSS Transition Management (802.11v) 551
Fast BSS Transition (802.11r) 559
WLAN SSID Profiles 560
WLAN Authentication 568

RF Planning and Channel Management 571
RF Management for Mobility Master Deployments with Managed Devices 571
RF Management for Deployments with a Stand-alone Controller 571
AirMatch RF Management Overview 572

ClientMatch Overview 574
Configuring AirMatch 577
Configuring ClientMatch 579
RF Management for Stand-alone Controller Deployments 580
802.11ad 586
ARM Coverage and Interference Metrics 586
Configuring ARM Profiles 588
Dynamic Bandwidth Switch 591
Troubleshooting ARM 591

Wireless Intrusion Prevention 593
Monitoring the Security Dashboard 593
Detecting Rogue APs 594
Working with Intrusion Detection 597
Configuring Intrusion Protection 608
Configuring the WLAN Management System 611
Understanding Client Blacklisting 618
Working with WIP Advanced Features 621

Access Points 627
Before Deploying an AP 627
Duplicate Address Detection 633
In the CLI 636
Important Points to Remember 637
Overview of Basic Functions of an AP 637
AP Configuration Profiles 639
Converting APs to Instant APs 642
Configuring Installed APs 643
Validating and Optimizing AP Connectivity 645
AP Groups 647
AP Image Preload 650
Enable and Configure AP Image Preload 651
View AP Preload Status 652
AP Discovery Logic 653
AP Channel Scanning 666
Managing AP Console Settings 668
Link Aggregation Support 672
On Mobility Master: 673
On an L2-connected High Availability (HA) standby or HA+VRRP controller: 673
On an L3-connected High Availability (HA) standby controller, or an L2- or L3-connected controller in dual-HA mode: 674
On Mobility Master 674
On an L2-connected High Avability (HA) controller that does NOT use dual-mode HA: 675
On L3-connected High Availability (HA) standby controllers, or HA controllers in dual HA mode, where each dual-mode controller acts as standby for the APs served by the other dual-mode controller: 675
2.4 GHz and 5 GHz Radio RF Management 676
High-Throughput APs 688
High-Efficiency (HE) APs 694
HE Pooling and Automatic Tri-Radio 696
Loop Protection 700
Support for Port Bounce 702
AP Packet Capture 703
Green AP 704

Air Slice 706
Optional AP Configuration Settings 708
AP USB Management 732
Disable AP Factory Reset 736

Secure Enterprise Mesh 738
Overview of Mesh Access Points 738
Overview of Mesh Links 740
Overview of Mesh Profiles 742
Overview of Remote Mesh Portals 746
Overview of AP Boot Sequence 747
Mesh Deployment Planning 748
Mesh Deployment Solutions 749
Mesh Configuration Procedures 752
Creating and Editing Mesh Radio Profiles 752
Creating and Editing Mesh High-Throughput SSID Profiles 757
Configuring Mesh Cluster Profiles 763
Configuring Mesh Clusters Associated with AP Groups 766
Configuring Ethernet Ports for Mesh 770
Configuring a Mesh Access List 773
Provisioning Mesh Nodes 774
Radio Selection for Mesh Links 775
Verifying Your Mesh Network 776
Configuring Remote Mesh Portals 776

Increasing Network Uptime With Redundancy Services 780
Mobility Master Redundancy Methods 780
In the CLI 783
Verifying VRRP Configuration 784
Logging and Debugging 785
In the WebUI 786
In the CLI 786
Verifying Master Redundancy 787
In the WebUI 788
In the CLI 788
AP and User Redundancy Methods 794
Active/Active Deployment Model 796
1:1 Active/Standby Deployment Model 796
N:1 Active/Standby Deployment Model 797

IP Mobility 809
Understanding Aruba Mobility Architecture 809
Configuring Mobility Domains 810
Tracking Mobile Users 812
Configuring Advanced Mobility Functions 814
Understanding Bridge Mode Mobility Deployments 822
Monitoring Network Traffic Using IP Flow Information Export 823
Enabling Mobility Multicast 826

External Firewall Configuration 831
Understanding Firewall Port Configuration in Aruba Devices 831
Enabling Network Access 832
Ports Used for VIA 832
Configuring Ports to Allow Other Traffic Types 833

Enhanced Security 834
Interoperability 834
Configuring PAPI Enhanced Security 834
Verifying PAPI Enhanced Security 835

Palo Alto Networks Firewall Integration 837
Pre-configuration on the PAN Firewall 837
Configuring PAN Firewall Integration 839

Remote Access Points 844
About Remote Access Points 844
Configuring the Secure Remote Access Point Service 845
Deploying a Branch or Home Office Solution 852
Bringing up Certificate-Based Remote AP in VMC 857
Enabling Remote AP Advanced Configuration Options 859
Adding or Editing a Remote AP Authorization Profile 873
Understanding Split Tunneling 873
Understanding Bridge 881
Provisioning Wi-Fi Multimedia 886
Reserving Uplink Bandwidth 886
Provisioning 4G USB Modems on Remote Access Points 887
Provisioning Remote AP at Home 889
Converting an Instant AP to Remote AP or Campus AP 892
Enabling Bandwidth Contract Support for Remote APs 893
Applying Contracts Per-Role 894
Applying Contracts Per-User 894

Virtual Intranet Access 897
License Requirements 897
Marking Outgoing Packets with ToS Bits 898
VIA Client Audit 898
VIA VPN Client Visibility 898
VIA VPN Client Capability 898
VIA Unique Identifier 899
VIA VPN Client Authentication 899

Spectrum Analysis 900
Understanding Spectrum Analysis 900
Creating Spectrum Monitors and Hybrid APs 905
Spectrum Analysis Tasks 907
Configuring Spectrum Analysis Dashboards 911
Customizing Spectrum Analysis Graphs 913
Working with Non-Wi-Fi Interferers 927
Understanding Spectrum Analysis Session Log 928
Viewing Spectrum Analysis Data 928
Recording Spectrum Analysis Data 929

Automatic Reporting (PhoneHome) 932
Registering with Activate 932
Configuring PhoneHome Automatic Reporting 933
Sending Reports to Activate vs. SMTP Servers 934
Sending an Individual Report 935
Viewing Report Status 935
PhoneHome-Lite 936

Management Access 938

Configuring Certificate Authentication for WebUI Access 938
Secure Shell 939
Enabling RADIUS Server Authentication 942
Default password recovery user 948
Alternate password recovery user 948
Connecting to AirWave Server 949
Custom Certificate Support for Remote AP 953
Implementing Specific Management Password Policy 955
Configuring Centralized Image Upgrades 957
Managing Certificates 972
Certificate Enrollment Using EST 979
Activate an EST profile using the CLI 982
Configuring SNMP 983
MIB Files 985
Enabling Capacity Alerts 985
Configuring Logging 986
Syslog Files 989
Enabling Guest Provisioning 989
Username and Password Authentication Method 994
Static Authentication Method 994
Smart Card Authentication Method 995
Username and Password Method 995
Static Authentication Method 995
Smart Card Authentication Method 995
Creating Multiple Guest Entries in CSV File 998
Importing CSV File into Database 999
Printing Guest Account Information 1002
Managing Files on Managed Device 1004
SCP Server Support 1008
Setting System Clock 1010
ClearPass Policy Manager Profiling with IF-MAP 1013
Whitelist Synchronization 1014
Downloadable Regulatory Table 1014
Infrastructure for Supporting Database Upgrade 1016
Configuring Concurrent Sessions 1017
Implementing Management User Audits 1017
Implementing Password Validation 1017
Maintaining Standard Mandatory Notice and Consent Banner 1018
Zeroizing TPM Keys 1018

Hotspot 2.0 1020

Hotspot Profile Configuration Tasks 1020
Access Network Query Protocol 1020
Hotspot Profile Types 1021
Configuring Hotspot 2.0 Profiles 1022
Configuring Hotspot Advertisement Profiles 1026
Configuring ANQP Venue Name Profiles 1028
Configuring ANQP Network Authentication Profiles 1030
Configuring ANQP Domain Name Profiles 1031
Configuring ANQP IP Address Availability Profiles 1031
Configuring ANQP NAI Realm Profiles 1032
Configuring ANQP Roaming Consortium Profiles 1034
Configuring ANQP 3GPP Cellular Network Profiles 1035

Configuring H2QP Connection Capability Profiles 1036
Configuring H2QP Operator Friendly Name Profiles 1037
Configuring H2QP Operating Class Indication Profiles 1038
Configuring H2QP WAN Metrics Profiles 1038

SDN Controller 1041
Southbound Interface 1041
SDN Controller Configuration on Mobility Master 1042
SDN Platform Services 1042
Northbound API 1051

OpenFlow Agent 1064
Enabling SDN Controller on Mobility Master 1064
Configuring OpenFlow Agent on Managed devices 1065
Viewing OpenFlow Information 1067

Loadable Service Module 1068
Service Modules 1068
Service Packages 1068
Upgrading a Service Module 1068
Troubleshooting 1070

Voice and Video 1072
Voice and Video License Requirements 1072
Configuring Voice and Video 1072
Working with QoS for Voice and Video 1083
UCC 1089
Understanding Extended Voice and Video Features 1124
WebRTC Prioritization 1131

IoT 1139
Payload Encoding 1140
Payload Content 1140
Transport Interval 1140
Filter Attributes 1140
Cell Size 1140
Threshold Attributes 1140
Out of Range Timeout 1140
Vendor Filter 1141
IoT and Third Party Server Integration 1141
Proxy Server 1141
Support for Zigbee USB Dongle 1141
Support for SES-Imagotag Cloud TLS Authentication 1142
Support for AmberBox Sensor 1142
Support for MySphera Tag 1142
Support for ABB Sensor 1142
Wi-Fi RTLS and BLE Telemetry Streams 1143
ZF Openmatics Support for ZF BLE Tag Communication 1143
Support for IoT Southbound API 1143
Aruba Sensor Value 1143
Reporting Sensor Values 1143
Favendo Sensors 1144
HanVit Sensors 1144
SoluM ESL Gateway 1145
Hanshow ESL Gateway 1146
Wiliot Sensors 1146
ZigBee Sockets 1147

Support for Input-Filter on BLE Devices ... 1147
Data Filter ... 1148
BLE Device Management ... 1150
Configuring IoT ... 1151

AirGroup ... 1161
Zero Configuration Networking ... 1161
AirGroup Solution ... 1162
AirGroup in ArubaOS ... 1162
AirGroup Services ... 1163
AirGroup Modes ... 1164
AirGroup Deployment Models ... 1164
Deprecated AirGroup Features ... 1165
AirGroup Features ... 1165
Prerequisites to Enable AirGroup ... 1172
Configuring AirGroup ... 1175
Best Practices and Limitations ... 1180
Troubleshooting and Log Messages ... 1183

Instant AP VPN Support ... 1186
Overview ... 1186
VPN Configuration ... 1191
Viewing Branch Status ... 1192
Backward Compatibility ... 1194

External Services Interface ... 1195
Sample ESI Topology ... 1195
Understanding the ESI Syslog Parser ... 1197
Configuring ESI ... 1199
Adding a new parser rule ... 1205
Deleting a syslog parser rule ... 1206
Editing an existing syslog parser rule ... 1206
Testing a Parser Rule ... 1207
Sample Route-Mode ESI Topology ... 1207
Sample NAT-mode ESI Topology ... 1213
Understanding BRE Syntax ... 1217

External User Management ... 1220
Overview ... 1220
How the ArubaOS XML API Works ... 1220
Configuring an XML Server ... 1220
Creating an XML Request ... 1221
XML Response ... 1224
Using the XML API Server ... 1227
Sample Scripts ... 1232

Behavior and Defaults ... 1239
Understanding Mode Support ... 1239
Understanding Basic System Defaults ... 1240
Understanding Default Management User Roles ... 1248
Understanding Default Open Ports ... 1251

DHCP with Vendor-Specific Options ... 1255
Configuring a Windows-Based DHCP Server ... 1255
Enabling DHCP Relay Agent Information Option (Option 82) ... 1257
Enabling Linux DHCP Servers ... 1261

802.1X Configuration for IAS and Windows Clients 1262
Configuring Microsoft IAS 1262
Configuring Management Authentication Using IAS 1264
Window XP Wireless Client Sample Configuration 1267

Glossary of Terms 1269

Revision History

The following table lists the revisions of this document.

Table 1: *Revision History*

Revision	Change Description
Revision 34	Removed the **AP Settings Triggering a Radio Restart** section from the **Access Points** chapter.
Revision 33	Updated the **Features not Supported in Each Forwarding Mode** table of the **Understanding Mode Support** section.
Revision 32	Updated the **Configuring the Session ACL Allowing Tunneling** section of the **Remote Access Points** chapter.
Revision 31	Updated the **RF Management (802.11a and 802.11g) Profiles** section of the **Secure Enterprise Mesh** chapter.
Revision 30	Updated the **AP High Availability Overview** section of the **Increasing Network Uptime With Redundancy Services** chapter.
Revision 29	Added the **Configuring AP Image Preload** topic in the **Management Access** chapter.
Revision 28	Updated the **Controller Discovery Using a DHCP Server** section of the **Access Points** chapter.
Revision 27	Updated the **Intelligent Power Management**, **Intelligent Thermal Management**, and **Intelligent Power and Temperature Monitoring** sections under **Access Points** chapter.
Revision 26	Added a note to the **Mobility Master Configuration Hierarchy** section.
Revision 25	Removed **RADIUS CoA** from the **Features not Supported in Each Forwarding Mode** table in the **Behavior and Defaults** chapter.
Revision 24	Updated the IPv6 address example in the **IPv6 Support** chapter.
Revision 23	The following changes were introduced: ■ Added **AP Hardware Offload** section under **Access Points** chapter. ■ Updated a note under **Native IPv6 Support** section of the **IPv6 Support** chapter.
Revision 22	Updated the **Upgrading Cluster** section of the **Controller Clustering** chapter.
Revision 21	The following changes were introduced: ■ Updated the **Configuring RADIUS Server Authentication with VSA** section of the **Management Access** chapter. ■ Updated the **Role-Based ACL** section of the **Roles and Policies** chapter.
Revision 20	The following changes were introduced: ■ Added a note for **Deny inter user traffic** parameter under **Virtual AP Profile Parameters** table of the **Configuring WLANs** chapter. ■ Added a note related to webserver configuration under **Managing Certificates** section of the **Management Access** chapter.
Revision 19	Updated the **Requirements** section of the **Controller Clustering** chapter.

Table 1: *Revision History*

Revision	Change Description
Revision 18	Updated the **Configuring Secondary Mobility Master Using Layer-3 Redundancy** section of the **Increasing Network Uptime With Redundancy Services** chapter.
Revision 17	▪ Updated the **Converting APs to Instant APs** section of the **Access Points** chapter. ▪ Updated **Configuring Additional RF Management Settings** section of the **Access Points** chapter. ▪ Updated the **Configuring VLANs** section.
Revision 16	The uplink load balancing feature is no longer supported.
Revision 15	The following changes were introduced: ▪ Added **Mesh Network with Mixed Indoor and Outdoor APs** section in the **Secure Enterprise Mesh** chapter. ▪ Updated the **Discovery Logic Workflow** section in the **Access Points** chapter.
Revision 14	Updated the **AP Packet Capture** section of the **Access Points** chapter.
Revision 13	Updated the **AP LACP Support** section of the **Controller Clustering** chapter.
Revision 12	Added **Hold Time** in VRRP parameters.
Revision 11	The following changes were introduced: ▪ Updated the **ArubaOS Multiversion Support** and **Enabling Linux DHCP Servers** sections. ▪ Updated the **Configuring a Windows-Based DHCP Server** section of the **DHCP with Vendor-Specific Options** chapter.
Revision 10	Updated the **Energy Efficient Ethernet** section of the **Access Points** chapter.
Revision 09	This following sections were updated: ▪ **Using ZTP with DHCP to Provision a Managed Device** section of the **Managed Devices at Branch Offices** chapter. ▪ **Sample Scripts** section of the **External User Management** chapter. ▪ **Configuring the AP System Profile** section of the **Access Points** chapter.
Revision 08	The following changes were introduced: ▪ Updated the description of interference index in the **ARM Coverage and Interference Metrics** section of **RF Planning and Channel Management** chapter. ▪ Updated the **Important Points to Note** section of the **Multiversion Support** chapter.
Revision 07	Updated a note in the **Obtaining Server Certificate** section of the **Management Access** chapter.
Revision 06	Updated the **Additional Video Configurations** section of the **Voice and Video** chapter.
Revision 05	Added a note to the **Replacing a Controller** section of **The Basic User-Centric Networks** chapter.

Table 1: *Revision History*

Revision	Change Description
Revision 04	Updated the **Infrastructure** section of the **Dashboard Monitoring** chapter.
Revision 03	Removed number of hops and configuring AirGroup server-based policy under AirGroup.
Revision 02	Removed references to Fremont devices.
Revision 01	Initial release.

This User Guide describes the features supported in ArubaOS 8.x and provides instructions and examples to configure Mobility Master, managed devices, and access points. This guide is intended for system administrators responsible for configuring and maintaining wireless networks and assumes administrator knowledge in Layer 2 and Layer 3 networking technologies.



Throughout this document, branch controller and local controller are termed as a managed device.

This chapter covers the following topics:

- New Features and Hardware Platforms on page 18
- Fundamentals on page 24
- System Requirements on page 27
- Supported Browsers on page 27
- Related Documents on page 28
- Conventions on page 28
- Contacting Support on page 29

What's New In ArubaOS 8.7.0.0

New Features and Hardware Platforms

This section lists the new features, enhancements, or hardware platforms introduced in ArubaOS 8.7.0.0.

Table 2: *New Features in ArubaOS 8.7.0.0*

Enhancements	Description
Air Slice	Aruba's key RF differentiation, Air Slice, designed for 11ax APs optimizes user experience and assures QoA to enterprise applications. Air Slice combines AppRF and UCC for classifying applications and it also supports custom flow definitions.
AP Boot Time	The AP fast boot is supported on AP-534, AP-535, and AP-555 access points for booting up within a minute with full functionality and configurations.
AP Fast Recovery	The 530 Series and 550 Series access points now support AP fast recovery feature for detecting firmware assert.
AP Master Discovery	ArubaOS now allows users to configure the preferred IP protocol for AP master discovery.

Enhancements	Description
AP USB Management	ArubaOS supports a new infrastructure to manage any USB device that is plugged to an AP. The infrastructure allows describing a USB device through either CLI configuration or by using predefined descriptors, USB device management through ACLs, plugin for USB devices, and sending notification to other processes.
BLE Device Management	ArubaOS now supports BLE device management through the IoT manager. The IoT manager allows configuration of BLE devices, provides visibility of detected BLE devices, and categorizes the BLE devices.
BLE Operation Mode	ArubaOS does not support BLE in 203H Series (AP-203H), 203R Series (AP-203R and AP-203RP), 207 Series (AP-207), 220 Series (AP-224 and AP-225), and AP-228 access points.
Captive Portal Authentication in Bridge Mode	ArubaOS supports captive portal authentication for VAPs in the bridge forwarding mode.
Enhancements to Centralized Licensing Feature	ArubaOS allows L3 redundant Mobility Masters to be configured as relay servers. For more details, see *ArubaOS 8.7.0.0 Licensing Guide.*
Configuring the Mobility Master IPv6 Address	It is now optional to configure the IPv4 address of the Mobility Master or the managed device in a native IPv6 deployment.
WebUI Support for Configuring IoT	ArubaOS allows configuration of IoT transport profile, IoT radio profile, ZigBee service profile, and ZigBee socket device profile through the WebUI. ArubaOS supports a new IoT device class filter **Exposure Notification**.
Cluster Dashboard Enhancements	Cluster dashboard is now visible in nodes other than the Managed Network.
Cluster Heartbeats Enhancements	New show commands are introduced to display more information about the cluster heartbeats.
Disable AP Factory Reset	ArubaOS allows users to disable AP factory reset while the AP is operational.
Dashboard Monitoring	The following changes are introduced: The **Dashboard > Overview** page displays the details of wired clients connected in bridge mode. The **Wireless Clients** table under **Dashboard** > **Overview** page displays all the IP addresses that are associated with a particular client.
Duplicate Address Detection (DAD)	The Duplicate Address Detection (DAD) feature identifies and prevents IP conflicts in an IPv6 deployment and ensures that the configured unicast IPv6 address is unique before it is assigned to a VLAN interface on the AP.
HE Pooling and Automatic Tri-Radio	ArubaOS supports HE dedicated radios, pooling of HE clients to HE-preferred radio, and automatic tri-radio mode.
IPv6 Support for RAPs in a Cluster	ArubaOS now assigns inner IPv6 pool to Remote APs to establish tunnel in a cluster deployment.
Mesh Links Table	The **Dashboard > Overview > Radios > Mesh Links** page is displayed with the following information: ▪ Cluster name

Enhancements	Description
	▪ Status ▪ Uplink Age ▪ Portal Access Point ▪ Mesh Points ▪ Band
Mesh Access List	The mesh access list feature allows each AP to establish mesh links only with the whitelisted neighboring APs.
Default Ciphers	The managed devices and APs use SHA2 and Group 14 as default values.
Fast BSS Transition (802.11r)	Fast BSS transition is now operational with WPA3-Personal, WPA3-Enterprise Basic, and WPA3-Enterprise Non-CNSA mode with GCM-256 encryption.
Green AP Support	AP-514, AP-515, AP-534, AP-535, AP-555, AP-504, AP-505, AP-505H, AP-518, and 570 Series access points support the Green AP, a power saving feature.
Hotspot 2.0 Support	530 Series, AP-555, and 570 Series access points now support Hotspot 2.0.
Hybrid Model Support	ArubaOS supports hybrid model that addresses issues regarding Remote AP termination on VMC.
Intelligent Temperature Monitoring System	ArubaOS supports the intelligent temperature monitoring system. When enabled, the temperature is dynamically controlled and the AP is allowed to cool down.
IoT Data Filter	ArubaOS supports data filters in the Iot transport profile that reduce the traffic on the telemetry interfaces.
IP Classification-based Firewall	ArubaOS supports IP classification-based firewall that help to identify the IP address and geolocation from where malicious activities originate.
IPv6 Support for CPPM Downloadable User Role	ArubaOS now allows to download the ClearPass Policy Manager (CPPM) user role using the IPv6 address.
Native IPv6 Support	ArubaOS provides native IPv6 support that allows enterprises to deploy pure IPv6 wireless network in a Mobility Master-Managed Device topology.
OpenSSL Upgrade	OpenSSL is upgraded from 1.01c to 1.02t for ArubaOS 8.7.0.0-FIPS.
Optimizing cluster live upgrade	The live upgrade process is optimized to take less time to upgrade all the devices in a cluster.
Radio Selection for Mesh Links	ArubaOS now allows to configure the 5 GHz radio used for mesh links. This feature is supported in dual-5 GHz and split-5 GHz radio enabled APs. This feature is designed to offer better control of the RF environment in mesh networks.
Multiple Policy Domains Support	ArubaOS now supports mulitple policy domains for group profiles.

Enhancements	Description
Spectrum Analysis Support	AP-504, AP-505, AP-505H, AP-514, AP-515, AP-518, and 570 Series access points now support spectrum analysis feature.
Support for Diffie-Hellman Groups in Enhanced Open Security	ArubaOS supports Diffie-Hellman Groups 20 and 21 for Enhanced Open security. Previously ArubaOS supported only Diffie-Hellman Group 19 for Enhanced Open Security.
Support for Input-Filter on BLE devices	ArubaOS supports a input-filter for BLE devices. Only BLE devices that should be reported are stored in the BLE-table.
Support for IPv4 and IPv6 APs in a Cluster	ArubaOS allows both IPv4 and IPv6 APs to connect to a cluster seamlessly, irrespective of the cluster IP address family.
Support for IoT Southbound API	ArubaOS supports an IoT Southbound API that allows interaction with IoT devices.
Support for New Modem	AP-303H Remote Access Point now supports GT Netstick GLU-194ST USB modem.
Support for SHA2 and DH Group	ArubaOS supports HMA-SHA2 and Diffie–Hellman Group 14 as the default Ciphers. Therefore, the managed devices and APs use SHA2 and Group 14 as default values. Also, a new crypto isakmp policy, 10015 is introduced to support SHA2 and Group 14.
Support for SHA-2 Authentication	ArubaOSsupports SHA-2 authentication for enhanced security and encryption.
Support for SSH Protocol on APs	ArubaOS allows SSH protocol over telnet protocol for high end encryption and enhanced security to avoid any network attack or risk of malicious users.
Support for Wiliot Sensor	ArubaOS supports Wiliot sensors that provide battery-free BLE tags.
Thermal Shutdown Support	ArubaOS supports thermal shutdown feature for all 802.11ax APs and for mesh mode APs.
Virtual Intranet Access	ArubaOS allows you to configure IPv6 address of the managed device in VIA connection profile.
VLAN Pooling Resiliency	The VLAN pooling resiliency feature automatically assigns clients to the next available VLAN ID if a particular VLAN pool is full.
Wi-Fi Uplink Support on 802.11r APs	ArubaOS now provides support for Wi-Fi uplink on all 802.11r enabled APs.
Web Server TLS version support	ArubaOS supports TLS v1.2 as the default SSL protocol in the Web Server.
Working with IPv6 RAs	ArubaOS provides support for DNS Search List (DNS-SL) option through IPv6 Router Advertisements that allows the IPv6 clients to resolve incomplete domain names.
ZigBee Sockets	Zigbee Socket Device may be configured and applied as a filter in IoT transport profile. With ZSD, specify the source endpoint, destination endpoint, destination profile ID or destination cluster ID and the packets between the ZigBee devices and server are transmitted through the Aruba Telemetry Websocket.

Table 3: *New Hardware Platforms in ArubaOS 8.7.0.0*



Check with your local Aruba sales representative on new managed devices and access points availability in your country.

Hardware	Description
9000 Series controllers	The 9012 controller is a wireless LAN controller that connects, controls, and intelligently integrates wireless Access Points (APs) and Air Monitors (AMs) into a wired LAN system. The controller has advanced IDS functionality and mobility services that is integrated with per user based enforcement policies for better security. The controller has an in build Bluetooth functionality and hardware integrated NFC support, with the following port configurations: ■ 12 x 10/100/1000BASE-T ports • 6 x PoP+ ports • MDI/MDX ■ 2 x USB 3.0 port ■ 1 x RJ45 console port ■ 1 x Micro USB console port For complete technical details and installation instructions, see *Aruba 9012 Controller Installation Guide*.
AP-505H Access Points	The Aruba AP-505H access points are entry-level, dual-radio wireless AP that can be deployed in either controller-based (ArubaOS) or controller-less (Aruba Instant) network environments. AP-505H delivers high performance concurrent 2.4 GHz and 5 GHz 802.11ax Wi-Fi (Wi-Fi 6) functionality with 2x2 MU-MIMO radios, while also supporting 802.11a, 802.11b, 802.11g, 802.11n, and 802.11ac wireless services. Additional features include: ■ IEEE 802.11a, IEEE 802.11b, IEEE 802.11g, IEEE 802.11n, IEEE 802.11ac, and IEEE 802.11ax operation as a wireless access point. ■ Compatible with IEEE 802.3bt, IEEE 802.3at, and IEEE 802.3af PoE standards. ■ One uplink Ethernet port capable of data rates up to 2.5 Gbps. ■ Four downlink Ethernet ports capable of data rates up to 1 Gbps, including two 802.3at PoE PSE ports for supplying power to downlink devices. ■ Integrated BLE and Zigbee radios. ■ Mesh ■ Flexible USB host interface with 5W power sourcing capability. For complete technical details and installation instructions, see *Aruba AP-505H Access Point Installation Guide.*
AP-518 Access Points	The Aruba AP-518 access points are high performance, multi-radio, outdoor access point that can be deployed in either controller-based (ArubaOS) or controller-less (ArubaInstant) network environments. These APs deliver high performance concurrent 2.4 GHz and 5 GHz 802.11ax Wi-Fi (Wi-Fi 6) functionality with MIMO radios (2x2 in 2.4 GHz, 4x4 in 5 GHz), while also supporting 802.11a, 802.11b, 802.11g, 802.11n, and 802.11ac wireless services. Additional features include:

Table 3: *New Hardware Platforms in ArubaOS 8.7.0.0*



Check with your local Aruba sales representative on new managed devices and access points availability in your country.

Hardware	Description
	■ IEEE 802.11a, IEEE 802.11b, IEEE 802.11g, IEEE 802.11n, IEEE 802.11ac, and IEEE 802.11ax operation as a wireless access point. ■ Two Ethernet ports, ENET0 and ENET1, capable of data rates up to 2.5 Gbps and 1 Gbps respectively. ■ Compatible with IEEE 802.3bt, IEEE 802.3at, and IEEE 802.3af PoE standards on both Ethernet ports. ■ Mesh ■ Thermal management For complete technical details and installation instructions, see *Aruba AP-518 Access Points Installation Guide.*
570 Series Access Points	The Aruba 570 Series access points (AP-574, AP-575, AP-577) are high performance, multi-radio, outdoor access points that can be deployed in either controller-based ArubaOS) or controller-less (Aruba Instant) network environments. These APs deliver high performance concurrent 2.4 GHz and 5 GHz 802.11ax Wi-Fi (Wi-Fi 6) functionality with MIMO radios (2x2 in 2.4 GHz, 4x4 in 5 GHz), while also supporting 802.11a, 802.11b, 802.11g, 802.11n, and 802.11ac wireless services. Additional features include: ■ IEEE 802.11a, IEEE 802.11b, IEEE 802.11g, IEEE 802.11n, IEEE 802.11ac, and IEEE 802.11ax operation as a wireless access point. ■ Two Ethernet ports, ENET0 and ENET1, capable of data rates up to 2.5 Gbps and 1 Gbps respectively. ■ Compatible with IEEE 802.3bt, IEEE 802.3at, and IEEE 802.3af PoE standards on both Ethernet ports. ■ Mesh ■ Thermal management For complete technical details and installation instructions, see *Aruba 570 Series Access Points Installation Guide*.
570EX Series Access Points	The Aruba 570EX Series access points (AP-575EX and AP-577EX) are high performance, multi-radio access points suitable for harsh and hazardous outdoor locations that can be deployed in either controller-based (ArubaOS) or controller-less (Aruba Instant) network environments. These APs deliver high performance concurrent 2.4 GHz and 5 GHz 802.11ax Wi-Fi (Wi-Fi 6) functionality with MIMO radios (2x2 in 2.4 GHz, 4x4 in 5 GHz), while also supporting 802.11a, 802.11b, 802.11g, 802.11n, and 802.11ac wireless services. The APs provide the following functionality: ■ IEEE 802.11a, IEEE 802.11b, IEEE 802.11g, IEEE 802.11n, IEEE 802.11ac, and IEEE 802.11ax operation as a wireless access point. ■ Two Ethernet ports, ENET0 and ENET1, capable of data rates up to 2.5 Gbps and 1 Gbps respectively. ■ Compatible with IEEE 802.3bt, IEEE 802.3at, and IEEE 802.3af PoE

Table 3: *New Hardware Platforms in ArubaOS 8.7.0.0*



Check with your local Aruba sales representative on new managed devices and access points availability in your country.

Hardware	Description
	standards on both Ethernet ports. ■ Thermal management For complete technical details and installation instructions, see *Aruba 570EX Series Access Points Installation Guide*.

Deprecated APs

The following APs are no longer supported from ArubaOS 8.7.0.0 onwards:

Table 4: *Deprecated AP Models*

Access Points Series	Model Numbers
100 Series	AP-104, AP-105
103 Series	AP-103
110 Series	AP-114, AP-115
130 Series	AP-134, AP-135
170 Series	AP-175AC, AP-175AC-F1, AP-175DC, AP-175DC-F1, AP-175P, AP-175P-F1
RAP 3 Series	RAP-3WN, RAP-3WNP
RAP 100 Series	RAP-108, RAP-109
RAP 155 Series	RAP-155, RAP-155P

Fundamentals

Mobility Master can be accessed through three different interfaces for maximum visibility and functionality:

- WebUI on page 24
- CLI on page 25
- JSON APIs on page 27

WebUI

Mobility Master supports up to 320 simultaneous WebUI connections. The WebUI is accessible through a standard Web browser from a remote management console or workstation. The WebUI includes configuration tasks. The tasks are:

- Provision New APs— Campus AP or Remote AP configuration.
- Create a New WLAN— Create and configure new WLAN(s) and associate with an AP group.
- Define WIP Policy— Define WIP policies and assign to AP groups.

- Bulk Configuration Upload— The Bulk Edit template (in Excel sheet) on the managed device allows you to specify the static IP assignment for individual managed devices.
- Upgrade Controllers— Upgrade the managed devices.
- Reboot Controllers— Reboot the managed devices.
- Show Upgrade Status— Display the upgrade status of the managed devices.

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

CLI

The CLI is a text-based interface accessible from a local console connected to the serial port on the Mobility Master or managed device or through a Telnet or SSH session.



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

- When entering commands remember that:
- commands are not case sensitive
- the space bar completes your partial keyword
- the backspace key erases your entry one letter at a time
- the question mark (?) lists available commands and options

Important Points to Remember

- The Mobility Master architecture spawns a new CLI session every time a user logs in to the CLI through Telnet, SSH, or Console. Since each CLI session is processed independently, multiple sessions do not block one another.
- See the ArubaOS *CLI Reference Guide* for more information on the new commands and parameters that are introduced to support new functions.
- Configurations must be performed in the context of a node in the configuration hierarchy. Users with the necessary privileges can change the node context on the CLI prompt.
- Users are required to commit configurations on Mobility Master before the configurations can be pushed and applied to the device.

Remote Telnet or SSH Session from Mobility Master

An administrator can initiate a remote telnet or SSH session from the Mobility Master to a remote host. The host can be a Mobility Master, managed device, or a non-Aruba host.



This feature is supported from the SSH session of the Mobility Master.

To initiate a telnet session from the Mobility Master to a remote host:

1. Initiate an SSH session to the Mobility Master.
2. Execute the **telnet <host> [port <port-num>]** command.

host: IPv4 or IPv6 address of the remote host.

port <port-num>: Telnet port number of the remote host. This is an optional parameter.

1. Once successfully connected, the remote host prompts the credentials. Enter the remote host credentials.

To initiate an SSH session from the Mobility Master to a remote host:

1. Initiate an SSH session to the Mobility Master.
2. Execute the **ssh <username> <ip_addr>** command.

username: Username of the remote host.

<ip-addr>: IPv4 or IPv6 address of the remote host.

Once successfully connected, the remote host prompts the credentials.

3. Enter the remote host credentials.

To end the remote host session, execute the **exit** command. The remote host displays the following message:

```
(host) [remote] #exit
Connection closed by foreign host.
(host) [mynode]#
```

Important Points to Remember

- The Mobility Master architecture spawns a new CLI session every time a user logs in to the CLI through Telnet, SSH, or Console. Since each CLI session is processed independently, multiple sessions do not block one another.
- New commands and parameters have been added to support new functions and provide increased visibility. See the ArubaOS *CLI Reference Guide* for more details.
- Configurations must be performed in the context of a node in the configuration hierarchy. Users with the necessary privileges can change the node context on the CLI prompt.
- Users are required to commit configurations on Mobility Master before the configurations can be pushed and applied to the device.

Limitations

This feature has few limitations. They are:

- This feature is supported from the SSH session of only the Mobility Master.
- There is an inactivity timeout for the CLI sessions. When an administrator initiates a remote session (inner) from the Mobility Master's SSH session (outer), and the remote session takes more time than the inactivity timeout session, the outer session times out although the inner session is active. The administrator has to log back in to the outer session once logged off from the inner session.
- Designated telnet client control keys do not work for remote telnet sessions. When an administrator initiates a remote telnet session (inner) from the Mobility Master's SSH session (outer), the designated telnet client control keys functions for the outer SSH session only. The administrator should designate unique control keys for each remote telnet sessions.

Seamless Logon

The Seamless Logon feature enables you to login from the Mobility Master to a managed device without entering a password. The user can remotely login from a centralized location (Mobility Master) to any managed device and execute the show and action commands. To login to a managed device, execute the **logon <device-ip>** command on the Mobility Master CLI:

```
(host) [mynode] #logon 192.0.2.22
Last login: Tue Jul 12 04:34:37 2016 from 192.0.2.81
(host-md) #
```

JSON APIs

JSON APIs are exposed for all configuration objects in Mobility Master and client location information from the ALE. Configuration APIs allow users to send configurations to Mobility Master and view those modifications through their own management system (CLI or WebUI). APIs in an operational state are also exposed. ALE APIs return client location information through the ALE server. Though most of this data is structured in the JSON format, some data may be arranged in a pre-formatted string. For more details on JSON APIs, refer to the *ArubaOS NBAPI Guide*. For more information about ALE APIs, refer to the *Analytics and Location Engine API Guide*.

System Requirements

Listed below are the minimum Hypervisor host system requirements for ArubaOS to run as a guest VM and the resources required for the VM to be functional:



It is not recommend to over subscribe the processors, memory, and NIC ports on the VM.

Table 5: *System Requirements*

Host Requirements	Aruba Mobility Master Virtual Appliance	Virtual Mobility Controller
Quad-core Core i5 1.9 GHz CPUs or Faster (hyper-threading enabled)	Minimum 3 cores (6 hyper-threading cores)	Minimum 2 cores (4 hyper-threading cores)
Memory	16 GB	8 GB
Physical NIC ports **NOTE:** One NIC port is shared with the host management and the second is reserved for datapath.	2	2
Disk space	64 GB	32 GB

Other Specifications

The Mobility Master runs on a virtual machine that is deployed through an OVF/OVA file.

Prerequisites for deploying the Mobility Master:

- vSphere Client 5.1 or 5.5 must be installed on a Windows machine. Support for vSphere Web Client and vCenter is available on ESXi versions 6.0 and 6.5.
- vSphere Hypervisor 5.1, 5.5, 6.0, or 6.5 must be installed on the server.
- An OVF/OVA template must be accessible from the ESXi host.
- VMware Enterprise Plus license must be installed on the Hypervisor.

Supported Browsers

The following browsers are officially supported for use with the ArubaOS WebUI:

- Microsoft Internet Explorer 11 on Windows 7 and Windows 8
- Microsoft Edge (Microsoft Edge 38.14393.0.0 and Microsoft EdgeHTML 14.14393) on Windows 10
- Firefox 48 and higher on Windows 7, Windows 8, Windows 10 and macOS
- Apple Safari 8.0 or later on macOS
- Google Chrome

Related Documents

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

- *ArubaOS Release Notes*
- *ArubaOS Getting Started Guide*
- *ArubaOS User Guide*
- *ArubaOS CLI Reference Guide*
- *ArubaOS API Guide*
- *Aruba Mobility Master Licensing Guide*
- *Aruba Virtual Appliance Installation Guide*
- *Aruba Wireless Access Point Installation Guide*

Conventions

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

Table 6: *Typographical Conventions*

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

The following informational icons are used throughout this guide:



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



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



Indicates a risk of personal injury or death.

Contacting Support

Table 7: *Contact Information*

Main Site	arubanetworks.com
Support Site	https://asp.arubanetworks.com/
Airheads Social Forums and Knowledge Base	community.arubanetworks.com
North American Telephone	1-800-943-4526 (Toll Free) 1-408-754-1200
International Telephone	arubanetworks.com/support-services/contact-support/
Software Licensing Site	lms.arubanetworks.com
End-of-life Information	arubanetworks.com/support-services/end-of-life/
Security Incident Response Team	Site: arubanetworks.com/support-services/security-bulletins/ Email: aruba-sirt@hpe.com

Mobility Master (ArubaOS 8.x.x.x) uses a centralized, multi-tier architecture under a brand new UI that provides a clear separation between management, control, and forwarding functions. The entire configuration for both the Mobility Master and managed devices is set up from a centralized point, thereby simplifying and streamlining the configuration process. Mobility Master consolidates all-master, single master-multiple local, and multiple master-local deployments into a single deployment model.

Whereas, the architecture of ArubaOS 6.x and earlier versions consist of a flat configuration model that contains global and local configurations. The global configurations are applied to the master controller which propagates those to its local controllers. The local configurations are applied to the master or the local controller directly.

Mobility Master takes the place of a master controller in the network hierarchy. Mobility Master oversees controllers that are co-located (on-premises local controllers or off-campus branch office local controllers). All the controllers that connect to Mobility Master act as managed devices.



Starting from ArubaOS 8.7.0.0, the master controller mode deployment model is not supported on any controller platform.

This section provides details on the following topics:

- Understanding Configuration Hierarchy on page 30
- Centralized Configuration on page 32
- Configuration Validation on page 37
- Serviceability on page 37
- Mobility Master User Interface on page 39

Understanding Configuration Hierarchy

The Mobility Master hierarchy simplifies the configuration process by supporting multiple configurations for multiple deployments using a single Mobility Master. Configuration elements can be mapped to one or more end devices, such as a managed device or VPN concentrator. Common configurations across devices are extracted to a shared template, which merges with device-specific configurations to generate the configuration for an individual device.

Figure 1 *Example of the Configuration Hierarchy*

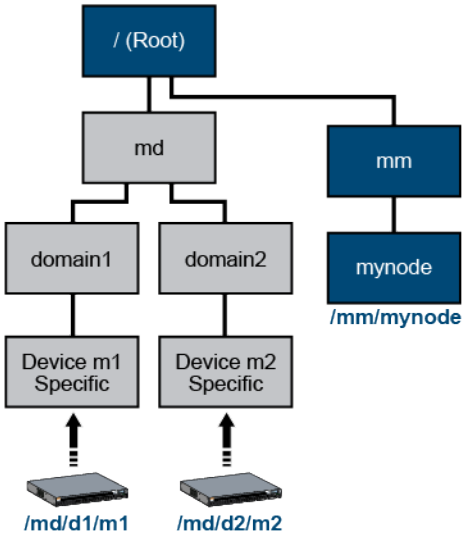


Figure 1 provides an example of the configuration hierarchy. The solid lines represent the hierarchy, the dotted arrows represent the device mapping, and each box represents a node in the hierarchy. When a device is added to Mobility Master, it must be mapped to a node or node-path in order to inherit configurations from the hierarchy. An explicit configuration node is also created for each device so that any device-specific configurations can be added directly to that node. Any device that is managed by Mobility Master is known as a managed device. For example, device **m2** in Figure 1 retrieves all device-specific configurations from the **Device m2 Specific** node. Since the **Device m2 Specific** node is mapped to the **domain2**, **md**, and **Root** nodes, the device also receives configurations from those nodes.

Each node contains a unique combination of common and device-specific configurations. The root node appears by default upon logging in to Mobility Master CLI.

The configuration hierarchy contains the following nodes and node structure:

Table 8: *Nodes and Node Structure*

Category	Node Name	Node Description
Mobility Master	/	Configurations common to Mobility Master and its managed devices (the root node). Configuration changes are not allowed on the root node.
	/md	Configurations common to all managed devices. The user can create additional nodes under this node.
	/mm	Configurations common to the primary and standby Mobility Master (VRRP pair).
	/mm/mynode	Configurations specific to a particular Mobility Master. This can only be edited on the respective Mobility Master.
Stand-alone Controller	/mm	Configurations common to the primary and standby stand-alone controllers (VRRP pair).
	/mm/mynode	Configurations specific to a particular stand-alone controller. This can only be edited on the respective stand-alone controller.



The term "mm" refers to Mobility Master and "md" refers to managed device.

Configurations for a node are obtained by traversing the node-path from the root node to the given node. For example, the **m1** device in Figure 1 receives configurations from all nodes along the **Root > md > domain1 > Device m1 Specific** node-path. Configurations that are set lower in the hierarchy (child node) can have more precedence than the same configurations set higher in the hierarchy (parent node), depending on the configuration type. In a single-instance configuration, such as the ESSID name, configurations from a child or device-specific node override common configurations from a parent node. In a multi-instance configuration, such as a server in an Auth Server group, configurations from a child node are placed in addition to the parent node configuration. For example, if a parent node specifies two radius servers and the child node specifies three radius servers, the device is provisioned with a total of five radius servers.

The configuration hierarchy is not the same as the physical topology. The hierarchy provides a simple way to organize configurations so that configuration elements can be shared across multiple devices without being duplicated. Configurations that are added to the root node, for example, are applied to all nodes within the hierarchy, while configurations that are only applied to a specific region override configurations for the corresponding child nodes. Order-dependent configurations, however, cannot be overridden. These configurations can only be set up once in the network hierarchy. For example, if a aaa server-group is configured on a parent node, it cannot be edited at the child node. Configuration hierarchies are tailored and organized to meet the unique needs of each customer.

Understanding the Node Hierarchy

You can view the hierarchy of the devices and groups on a Mobility Master at a global level. Mobility Masters are placed into the **/mm** group and managed devices are in the **/md** group.

/md—This is the global or root level where anything configured is applicable to all the nodes globally. It is recommended not to edit or add additional configuration at this level.

/md/<group name>— This is used to differentiate the sites physically or by the type of deployment such as DMZ, Branch, Campus, RAPs, and so on.

When you log in to the Mobility Master, you are placed in the **/mm/mynode** prompt by default.

Navigating through Node Hierarchy

You can use either the **change-config-node** or **cd** command to navigate to any node from the current node.

Both commands auto complete the group or folder names. You can also use the device hostname as an alias to navigate to a device node in the hierarchy. In doing so, your prompt changes to reflect where you are in the hierarchy:

```
(host) [mynode] #change-config-node Aruba7010
(host) [00:0b:86:99:97:57] #
```

The following CLI command displays your current node:

```
(host) [00:0b:86:99:97:57] #pwd
/md/Home-Production/00:0b:86:99:97:57
```

The following CLI command allows you to navigate one group up in the hierarchy:

```
(host) [00:0b:86:99:97:57] #cd ..
(host) [Home-Production] #
```

Centralized Configuration

Mobility Master uses a centralized configuration application to maintain all configurations under the management domain, eliminating the use of multiple points of contact to apply global and local

configurations to each managed device. You can organize all common configurations at a higher level of the hierarchy.

This section includes the following topics:

- Mobility Master Configuration on page 33
- Allowed Node Operations on page 33
- Access Permissions on page 34
- Bulk Edit on page 34
- Centralized Configuration on page 32
- Validation and Application Processes on page 36

Mobility Master Configuration

The Mobility Master that provides this configuration service to other devices in the network also contains its own configuration. The Mobility Master configuration is obtained through nodes in the hierarchy labeled **/mm** or **/mm/mynode**. Configurations under the **/mm** node, which are shared by the redundant Mobility Master pair (primary and standby Mobility Masters), are synced to the standby Mobility Master. Configurations under **/mm/mynode** are synced to individual Mobility Master devices.

Allowed Node Operations

The following node operations are allowed on Mobility Master:

- **Create Node**: Creates a new node as the child of an existing node in the configuration hierarchy (system-generated or user-created)
- **Add Device**: Associates a device to an existing node in the hierarchy. This device inherits configurations from all nodes between the root node and the device (node-path).
- **Delete Node**: Deletes an existing user-created node or node without any child nodes. System-generated nodes cannot be deleted. Only leaf nodes without any child nodes can be deleted.
- **Delete Device**: Deletes a currently associated device from the configuration hierarchy. This will cause the device to reload and erase all configurations received from Mobility Master.
- **Clone Node**: Copies the configuration of an existing node into a new node. The new node is created as a child of an existing node in the hierarchy.
- **Move Node**: Moves an existing user-created node in the hierarchy to the specified destination node. System-generated nodes cannot be moved. Ensure the following points while moving a node or device, otherwise the move operation will fail:
  - The node to be moved is a leaf node and does not have any group node or a device node as a child node under it.
  - No configuration is pending on the parent nodes of the child node to be moved.
  - The configuration on the node to be moved is complaint with the configuration in the new ancestor nodes chain.
- **Rename Node**: Renames the existing node name to the specified name. The node paths of the child nodes under the renamed node are automatically updated.
- **Drag and Drop Node**: Allows you to move any controller from one group to another group within the hierarchy, without deleting the controller from the Mobility Master.



Moving multiple controller or group within the network hierarchy is not supported.

- **Edit Action**: Allows you to rename a controller or a group in the managed network hierarchy.

Refer to the ArubaOS *Command Line Interface Reference Guide* for more details on the configuration commands for node and device management.

Access Permissions

The Mobility Master management domain can be large and widespread across various geographic regions. In a Mobility Master, the editing scope of the admin user can be restricted to individual node-paths within the configuration hierarchy, unlike the legacy ArubaOS management domain where an administrator can modify any configuration in the system.

Each management user is granted editing permissions for a given node, allowing the user to modify the configuration for that node and any child node within its node-path. The user, however, cannot modify any parent nodes or nodes on a different path in the hierarchy. Users can view configurations for any node in the hierarchy to refer to a parent node configuration or verify that the derived configuration for a device matches the parent node configuration.

- Management users that are configured with the root (**/**) or Mobility Master (**/mm**) nodes are granted editing permissions for Mobility Master.
- Management users that are configured with permissions to the mynode (**/mm/mynode**) can modify configurations under **/mm/mynode** for the respective Mobility Master and stand-alone controller.
- Management users that are configured with permissions to a managed device can modify configurations for that managed device.
- Only the management users that are configured with root node level permissions can modify configurations on both Mobility Master and managed devices.

Bulk Edit

The Bulk Edit Support feature enables you to perform a bulk configuration of managed devices in the Mobility Master. This option helps reduce the time taken to perform configuration tasks individually.

The following procedure describes how to do a bulk edit:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Tasks > Bulk configuration upload.**
2. Click **Download** sample template.
3. Enter values in the fields provided in the template.
4. Save the file.
5. Select **Browse** and navigate to the path where the template is stored.
6. Click **Submit**.

The **Bulk Configuration Status** pop up is displayed with the status of the configurations that are being applied. Once the configurations are applied successfully, a message confirming that the file upload was successful is displayed. The next pop up displays the following details:

- **Timestamp**
- **Status**
- **Number of devices updated**
- **Total new devices added**

If the configurations are not applied successfully, the **Bulk Configuration Status** pop up displays the reason for the failure and the managed device will rollback to the previous configuration.



When devices are added using the Bulk Edit feature, each template file can include up to 400 devices.

ArubaOS now provides IPv6 support for the Bulk Edit feature. Hence, the Bulk Edit template file now contains all the IPv6 commands that are required for native IPv6 deployment..

Bulk Stand-alone Controller Deployment

This feature supports bulk configuration of stand-alone controllers by replacing the configuration files in the stand-alone controllers and rebooting them.

The following procedure describes how to replace the configuration files in the stand-alone controller:

1. Login to the node, **/mm** or **/mm/mynode** from which you want to copy the configuration files and execute the command, **encrypt disable**.
2. Execute the command, **show configuration committed** and save the configuration to a .cfg file (for example, mmfile.cfg) on the tftp server.
3. Edit the file and copy that file to the controller's flash from the server using the command,

   (host) [mynode] # copy tftp: <IP address> vmccfg1.cfg flash: vmccfg1.cfg
4. Execute the command, **replace-config-reboot**, configuration node replace-config-reboot <filename1> <config-path1> <filename2> <config-path2>

   For example, configuration node replace-config- reboot mmfile.cfg /mm mynodefile.cfg /mm/mynode

   The controller prompts you to upload both the **/mm** and **/mm/mynode** files together.
5. Once the command is executed, the stand-alone controller will prompt you to reboot the controller.
6. Reboot the controller and the stand-alone controller will now boot up with the copied configuration.

Override Support in the WebUI

Starting from ArubaOS 8.2.0.0, the Mobility Master WebUI provides an option to retain or remove overrides for the fields configured under a node. If any field has an overridden value, the UI displays a blue dot to the left of the field name. Clicking on the dot, gives you an option to remove the overrides.



Order-dependent configurations, such as roles and ACLs, cannot be overridden. These configurations can only be set up once in the network hierarchy.

Support for Viewing Inheritance History in the WebUI

Starting from ArubaOS 8.3.0.0, the WebUI allows you to view the inheritance details of any configuration at any group or node level. This feature is supported only for configurations that can be overridden. A blue color information icon is displayed in the respective rows of the configuration table under which some configurations are overridden. Clicking on the icon displays the details of the inheritance with a link to the parent node. You can click on the parent node link to navigate to the parent node level. You can choose to remove all the overrides under the selected node level from this pop-up window by

clicking the **Remove Overrides** button. Else, you can choose to remove the individual configuration overrides at the field level.

Validation and Application Processes

When a user enters a configuration into a managed device, the configuration is validated. The validated configuration is accepted by the system but does not take effect until the configuration is committed. When the configuration is being committed, it is stored in the persistent memory, allowing users to verify the configuration before making it operational.

This separation of validation and application processes is applied to both the Mobility Master and managed devices. Since each node can be managed by a different admin user, the commit operation is executed on a per-node basis and follows the configuration hierarchy. For example, if a configuration has a dependency, the dependent configuration must be present on that node or one of the parent nodes .

Configurations are classified as pending configuration or committed configuration. A pending configuration refers to a configuration that has been validated but not yet committed. A committed configuration refers to all configurations that have been committed by the user. Users can view pending configurations at any time to commit, purge, or leave the configuration uncommitted. Pending configurations are only allowed on one node at any given time in a given configuration sub-tree.

Viewing the Node Hierarchy

The following CLI command displays how the devices and groups are organized at a global level:

```
(host) [mynode] #show configuration node-hierarchy
```

The following sample output displays the list of devices and nodes configured under the root node.

```
Default-node is "/md". Autopark is enabled.
Configuration node hierarchy
----------------------------
Config Node                 Type    Name
-----------                 ----    ----
/                           System
/md                         System
/md/00:0b:86:99:e2:17       Device
/md/VPNC                    Group
/md/VPNC/00:1a:1e:01:46:38  Device
/md/VPNC/00:1a:1e:02:03:d0  Device
/mm                         System
/mm/mynode                  System
```



The **show running-config** command from the Mobility Master displays the configuration on the Mobility Master and not on the other nodes or managed devices.

Viewing Configuration on Nodes

Use the following variants of the show commands to view the configuration information on a node or device level:

- **show configuration effective**—Displays the running configuration of the current node. You can also view the configuration on a specific node from a different node by specifying the absolute path of the node in the command.
- **show configuration effective detail**—Displays the full configuration details on your current node. It also indicates if a configuration is inherited from a group level or local to the managed device.

- **show configuration committed**—Displays the configuration that is only local to a specific node and not inherited from a parent node in the hierarchy. Configurations such as IP addresses and hostnames are some examples.
- **show configuration pending**—Displays the configuration details which are yet to be committed to the managed device or group, that is any configuration changes that are made before executing the **write memory** command or submitting the pending changes in the WebUI. This is used to review any configuration before it is applied from the Mobility Master to the managed devices. The output of the command is relevant only to the current node.
- **show configuration partial**—Displays the incremental change in the configuration between the last two synchronizations from the Mobility Master to the node.
- **show configuration similar**—Displays the like configuration between two specific nodes or devices. This is useful to verify equal settings between groups or devices. The output displays only the configurations that are same between both nodes. If you are comparing devices, you must use the path as displayed in the output of the **show configuration node-hierarchy** command.
- **show configuration diff**—Displays the configurations that are different between two specific nodes or devices. A minus sign against a configuration indicates that it is present in the node specified first but absent in the second node. A plus sign indicates that the configuration is absent in the first node but present in the second node.

For more information on various configuration show commands, *see ArubaOS 8.x CLI Reference Guide*.

Configuration Validation

Mobility Master uses a centralized validation model that performs various types of validations for different targets. Configuration validation falls under one of the following categories:

- **Syntax Validation**: Basic parser validations (for example, making sure the syntax of a command is correct, the data type is correct, or a value is within a valid range).



- Roles, ACLs, and pools (DHCP, VLAN, tunnel, and NAT) must be written in lower-case. Passwords, crypto keys, and ESSIDs can be written in both upper-case and lower-case.

- **Semantic Validation**: Custom application-specific validations (for example, dependency checks across commands or instance count limits). Dependency checks are limited to the nodes from which the target device inherits the configuration.
- **Platform Validation**: Platform model-specific validations (for example, determining which features are supported on a platform or the type and count of ports on a platform).



Validation is not available on the setup dialogue. Users must manually verify the setup dialogue information for each managed device.

Validation Failures

If a command does not pass validation, it is rejected and will not be included in the pending configuration for that node. If a new device that cannot support an existing configuration is added, the device add is rejected.

Serviceability

Managed devices are always serviceable from the centralized management location. When a managed device boots up for the first time under the factory default state, it auto-provisions and establishes connectivity to Mobility Master through ZTP. Managed devices can also be provisioned manually through the setup dialog box. Managed devices can encounter connectivity loss due to bad configurations, network connectivity issues, and so on. The system attempts to recover from these situations when possible.

This section includes the following topics:

- Bad Configuration Recovery on page 38
- Disaster Recovery on page 39
- Initial Provisioning Recovery on page 39

Bad Configuration Recovery

Certain configurations, such as those in the following list, can interfere with the connectivity between managed devices and Mobility Master:

- Uplink port shut
- Partially configured uplink VLAN
- Limiting bandwidth contract policy
- Bad ACL

Bad configurations can be caused by simple typo errors. Even if the user discovers the error, the bad configuration may have already caused connectivity loss, preventing the user from pushing the correct configuration to the managed device.

Mobility Master supports an auto-rollback mechanism that reverts the managed device to the last known good configuration prior to the management connectivity loss. Mobility Master also indicates if a device has recovered from a bad configuration through the **show switches** command output. The output for this command labels the **Configuration State** for the managed device as **CONFIG ROLLBACK** if the device has recovered connectivity using the rollback configuration. When the user fixes the bad configuration on Mobility Master, the managed device recovers automatically, and the state changes to UPDATE SUCCESSFUL.

Example output for the **show switches** command:

```
(host) [mynode] #show switches

Thu Jun 09 12:13:45.735 2016

All Switches
------------
IP Address     IPv6 Address  Name               Location          Type     Model
Version                  Status  Configuration State  Config Sync Time (sec)  Config ID
----------     ------------  ----               --------          ----     -----
-------                  ------  -------------------  ----------------------  ---------
192.192.192.1  None          TECHPUB_MASTER     Building1.floor1  master   ArubaMM
8.0.0.0-svcs-ctrl_55038  up      UPDATE SUCCESSFUL    0                       27
192.192.192.2  None          TECHPUB_STANDBY    Building1.floor1  standby  ArubaMM
8.0.0.0-svcs-ctrl_55038  up      UPDATE SUCCESSFUL    10                      27
192.192.189.1  None          TECHPUB_LC1_189.1  Building1.floor1  MD       Aruba7010
8.0.0.0-svcs-ctrl_55038  up      UPDATE SUCCESSFUL    0                       27
192.192.192.3  None          TECHPUB_x86_LC     Building1.floor1  MD       VMC-TACTICAL
8.0.0.0-svcs-ctrl_55038  up      UPDATE SUCCESSFUL    0                       27
192.192.189.2  None          TECHPUB_LC2_189.2  Building1.floor1  MD       Aruba7005
8.0.0.0-svcs-ctrl_55038  up      UPDATE SUCCESSFUL    0                       27
```

```
Total Switches:5
```

Disaster Recovery

If auto-rollback from a bad configuration fails, and connectivity between the managed device and Mobility Master remains disrupted, users can enable **Disaster Recovery** mode on the managed device using the **disaster-recovery on** command. Under the regular mode, the **/mm** node downloads configurations from Mobility Master that cannot be modified directly on each managed device. **Disaster Recovery** mode grants users access to the **/mm** node through the managed devices while blocking any further configuration synchronizations from Mobility Master. With full control of the **/mm** node, users can make local modifications on each managed device to restore connectivity to Mobility Master.



Local configurations are only used for debugging purposes and are not visible on the Mobility Master.

After connectivity is restored and verified, the user must fix the configuration on Mobility Master and exit the **Disaster Recovery** mode. When the user exits **Disaster Recovery** mode from the managed device, a full configuration sync is triggered between the managed devices and Mobility Master, which now contains the latest effective configurations.

The following CLI command enables **Disaster Recovery** mode:

```
(host-md) #disaster-recovery on
******************************
Entering disaster recovery mode
******************************
(DR-Mode) [mm] #
```

The following CLI command disables **Disaster Recovery** mode:

```
(DR-Mode) [mm] #disaster-recovery off
```

Initial Provisioning Recovery

If the managed devices fail to connect to Mobility Master on multiple attempts during the initial provisioning process (for example, when the Mobility Master IP or FQDN is entered incorrectly in Aruba Activate), the managed device deletes all provisioning information and restarts the auto-provisioning process. The user is expected to correct the provisioning information under Aruba Activate. After the provisioning information is corrected, the managed device automatically recovers during the next auto-provisioning attempt.

Mobility Master User Interface

The Mobility Master user interface provides ease-of-use through an intuitive layout and simple navigation model.

Navigation Model

Each page of the Mobility Master UI is divided into the following sections:

- **Header**, which includes the following:
- **Aruba logo**: The Aruba logo.
- **Deployment mode and hostname**: The deployment mode and hostname of the Mobility Master or managed device.

- **Network Status Counters**: Counters for reachable and unreachable controllers, reachable and unreachable access points, clients, and alerts.
- **Help**: Initiates help mode to display available help information in the UI. See Help Mode on page 42 for more details.
- **User menu**: Drop-down menu that displays your username. It allows you to logout of the Mobility Master or managed device. The **Preferences** option allows you to enable or disable the **Profiles** link in the following pages:
  - **All Profiles** table of the Mobility Master node.
  - **WLANs** table and **AP Group** table of the Managed Device node.



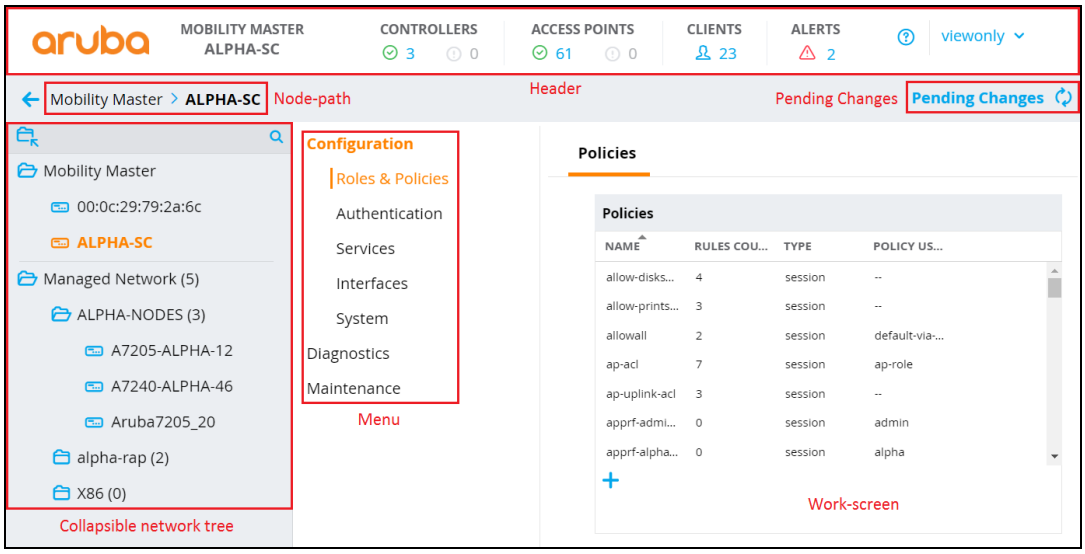
The **Profiles** link is displayed only when the **show advanced profiles** check box is selected in the **Preferences** option of the User menu.

Limitations

- Advanced profile configuration is controller specific (domain name)
- Advanced profile configuration is not per-user specific
- It is browser specific, irrespective of user login—for example, if a user enabled Preferences in the Chrome browser it will not carry forward to IE or Firefox.

- **Node-path**: Node-path within the network hierarchy.
- **Pending Changes**: List of all pending configuration changes. See Pending Changes on page 41 for more details.
- **Menu**: Main menu, which includes the **Dashboard**, **Configuration**, **Diagnostics**, and **Maintenance** menu items. Select a menu item to reveal the corresponding sub-menu items. See Navigation Levels on page 41 for more details.
- **Collapsible network tree**: Complete network hierarchy that is revealed or hidden when you click the menu or arrow button, respectively, next to the node-path. See Network Tree on page 40 for more details.
- **Work-screen**: Content description for a menu item or tab.

Figure 2 *Overview of the User Interface*



Network Tree

The Mobility Master UI allows users to create, modify, and delete any node in the network hierarchy from a central location. By clicking the menu button next to the node-path, you can reveal the entire network hierarchy. Select a node to further expand the hierarchy and display the corresponding child nodes. The network can be organized in a hierarchy of up to five levels, including groups, sub-groups, and the managed devices that are added to these groups.

When a node is selected in the network hierarchy, any configuration changes made from the UI are applied to the selected node and sub-tending managed devices.



If you are logged into a device that is managed by a Mobility Master or legacy master controller, the hostname of the Mobility Master or master controller is displayed in the node-path.

For more information on the configuration hierarchy, see Mobility Master Configuration Hierarchy on page 30.

Navigation Levels

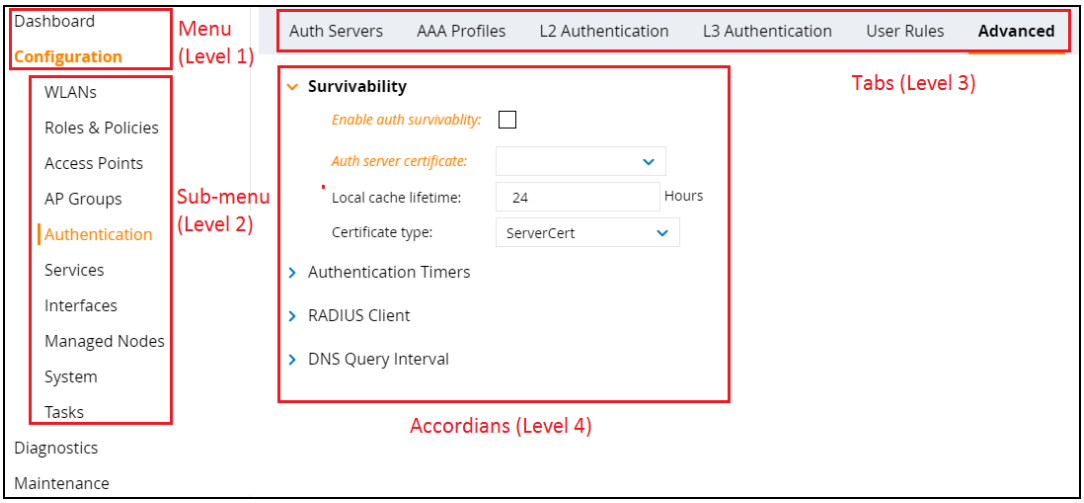
The Mobility Master navigation model is organized into four levels:

- **Level 1**: Menu (for example, Configuration, Diagnostics, and Maintenance)
- **Level 2**: Sub-menu (for example, Authentication, Interfaces, and Services)
- **Level 3**: Tabs (for example, Auth Servers, AAA Profiles, and Layer-2 Authentication)
- **Level 4**: Accordions (for example, Survivability and Authentication Timers)

Each Level-1 item can be expanded to display the corresponding Level-2 items. Each Level-2 item is further expanded to organize and group content on the work-screen. Based on the following dependencies, certain menu, tab, or accordion items may be visible or hidden in the UI:

- Selected node
- License
- Controller model
- User Role

Figure 3 *WebUI Menu, Tabs, and Accordions*



Pending Changes

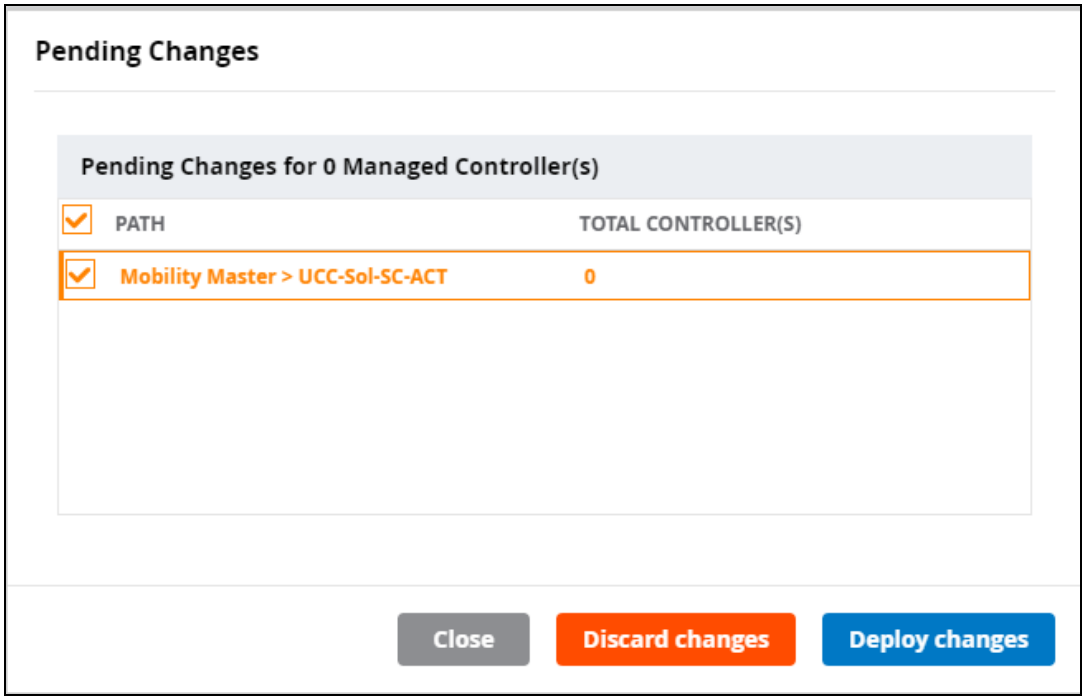
Commands are executed when a user clicks **Save** or **Submit**. The **Save** or **Submit** buttons are disabled by default and can only be enabled when the user modifies a configuration on the page. When a user

clicks the **Save** or **Submit** button, the configuration change is pushed to the **Pending Changes** zone of Mobility Master. Modifications are not applied to the network until all pending changes are deployed. Click **Pending Changes** to view, deploy, or discard all pending modifications.

NOTE

Nodes cannot be edited if any parent or child node contains undeployed pending changes.

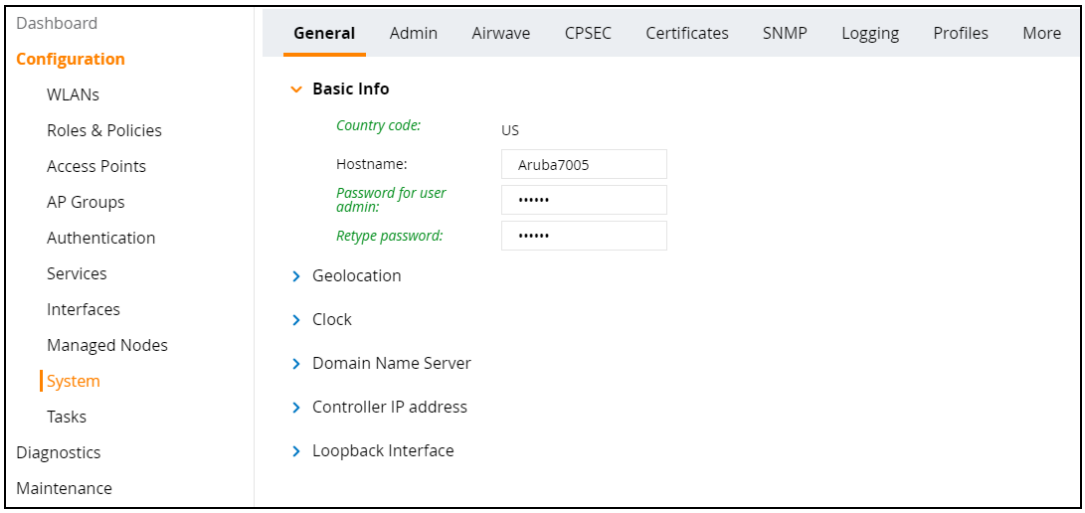
Figure 4 *Pending Changes Window*



Help Mode

The **Help** button in the header section of the UI allows you to switch the system to help mode. All non-active labels that appear in green italics indicates that help information is available for these labels. Mouseover or click any green italic label to view the help information for that field in a pop-up window.

Figure 5 *Help Mode in the WebUI*



Tables

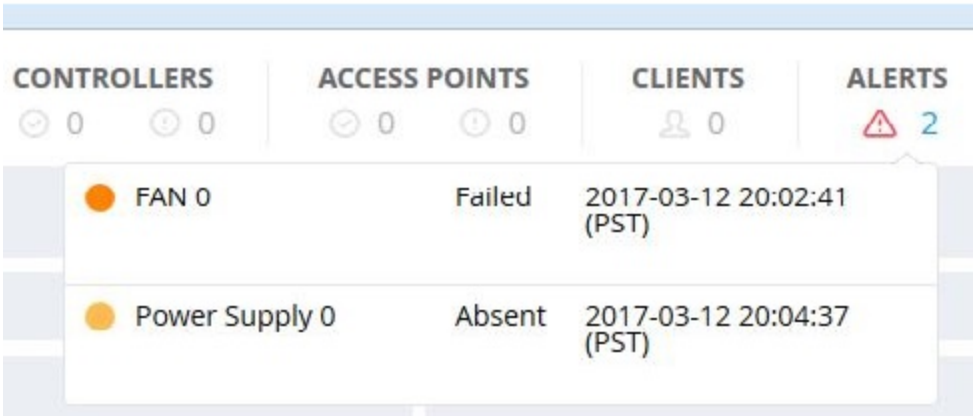
Mobility Master presents data and configuration information through two types of tables: primary tables and secondary tables. Primary tables display the main object of the page at the top of the screen (for example, the **Server Groups** table under **Configuration > Authentication > Auth Servers**). By selecting a row from the table, you can view and modify the configuration parameters for that entry in an editing pane that is displayed at the bottom of the screen.

The secondary table is located within the editing pane of a selected row and provides more in-depth information on each entry. For example, when you select the **default** server group entry from the **Server Groups** (primary) table, the secondary **Server Group > default** table is displayed at the bottom of the screen.

Alerts

Starting with ArubaOS 8.1.0.0, alerts for Mobility Master Hardware Appliance are enabled in the WebUI. Figure 6 shows an example of a Mobility Master Hardware Appliance alert.

Figure 6 *Mobility Master Hardware Appliance Alert in WebUI*



The Mobility Master Hardware Appliance alerts are as follows:

Table 9: *Mobility Master Hardware Appliance Alerts*

Component	State	Indicator	Status
Power Supply	Absent	Yellow	Minor Alarm
	Failed	Amber	Major Alarm
Temperature	Temperature more than 38°C	Yellow	Minor Alarm
	Temperature more than 40°C	Amber	Major Alarm
	Temperature more than 45°C	Red	Critical Alarm
Fan	Absent	Amber	Major Alarm
	Failed	Amber	Major Alarm

Starting with ArubaOS 8.2.0.0, Mobility Master provides the essential infrastructure for multiversion support across all managed devices in the network. With this enhancement, the ArubaOS version on each managed device can be different from that in the Mobility Master in the network.

The multiversion infrastructure performs the centralized validation for the configurations of different ArubaOS versions run on the managed devices. The configurations that are not compatible with the managed device's ArubaOS version will not be sent to the managed device.

This feature supports the following scenarios:

- Customers want to upgrade only the Mobility Master with the latest ArubaOS version to use centralized services.
- Customers want to upgrade only a few managed devices in their network with the latest ArubaOS version to test some features of their interest.
- Customers want to upgrade their network in certain geographical locations and plan to upgrade the entire network incrementally.

Important Points to Note

The following are important points to note before implementing the multiversion support in your network:

- ArubaOS 8.2.0.0 is the minimum supported version on the managed devices and the Mobility Master.
- The Mobility Master can run an ArubaOS version that is either the same or a higher version of ArubaOS than the versions on the managed devices; the minimum supported version on both platforms is ArubaOS 8.2.0.0.
- Multiversion is supported only if the Mobility Master is running up to two code versions higher than the code versions running on the managed devices. For example, multiversion is supported if a Mobility Master is running ArubaOS 8.5.0.0 and the managed devices are running either ArubaOS 8.3.0.0, ArubaOS 8.4.0.0, and/or ArubaOS 8.5.0.0. Managed devices running ArubaOS 8.2.0.0 or earlier versions and/or ArubaOS 8.6.0.0 or later versions are not supported.

UI Support for Multiversion

When the managed devices and the Mobility Master run different ArubaOS versions, the following rules apply:

- At all levels of hierarchy, the WebUI elements of the later ArubaOS version is always shown to the user.
- At the group level, the following rules apply:
  - All WebUI elements that are new in the ArubaOS version of the Mobility Master are shown.
  - The WebUI elements that are obsolete in the ArubaOS version of the Mobility Master are not shown.

- At the device level, the following rules apply:
  - The WebUI elements that are obsolete in the ArubaOS version of the Mobility Master but not obsolete on the ArubaOS version of the device are shown.
  - The WebUI elements that are obsolete on the ArubaOS version of the device are not shown.
  - The WebUI elements that are introduced in the ArubaOS version later than that on the device are not shown.

Display of ArubaOS Version Identifiers

The UI displays the ArubaOS versions running on the Mobility Master and the managed device:

The following changes are applicable in the WebUI of the managed device:

- **Mobility Master ArubaOS version identifier**: The ArubaOS version of the Mobility Master is displayed at the bottom of the left navigation pane. The Mobility Master version identifier is displayed as *Mobility Master: Version <version #>*.
- **Managed device ArubaOS version identifier:** If the ArubaOS version running on the managed device is different from that running on the Mobility Master, an information icon is displayed. It shows the ArubaOS version running on the managed device. The managed device version identifier is displayed as *Version <version #>*.



The managed device version identifier is not displayed when the ArubaOS version running on the managed device and the Mobility Master are the same.

- You can click the information icon to view the following details in the pop-up box:
  - The controller or managed device name and its ArubaOS version.
  - Mobility Master label and its ArubaOS version.

This chapter describes how to connect a managed device and an Aruba AP to your wired network. After completing the tasks described in this chapter, see Access Points on page 627 for information on configuring APs.

This chapter describes the following topics:

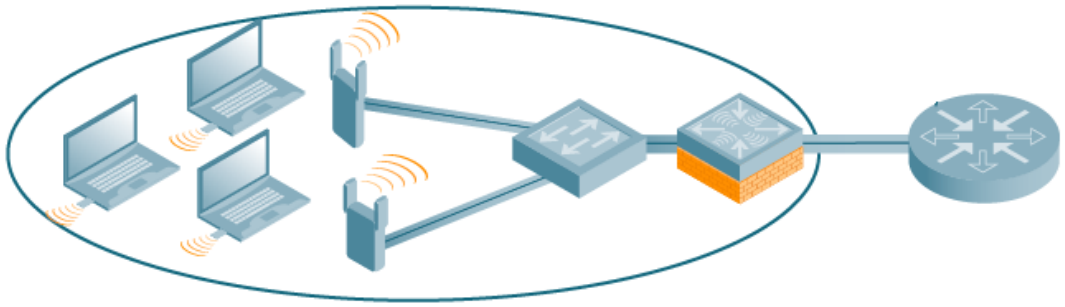
- Understanding Basic Deployment and Configuration Tasks on page 46
- Managed Devices Configuration Workflow on page 49
- Connect the Managed Device to the Network on page 50
- Using the LCD Screen on page 51
- Configuring a VLAN to Connect to the Network on page 54
- Configuring User-Centric Network on page 61
- Replacing a Controller on page 61

Understanding Basic Deployment and Configuration Tasks

This section describes typical deployment scenarios and the tasks you must perform while connecting to a managed device and Aruba AP to your wired network.

Deployment Scenario #1: Managed Device and APs on Same Subnet

Figure 7 *Managed Device and APs on Same Subnet*



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

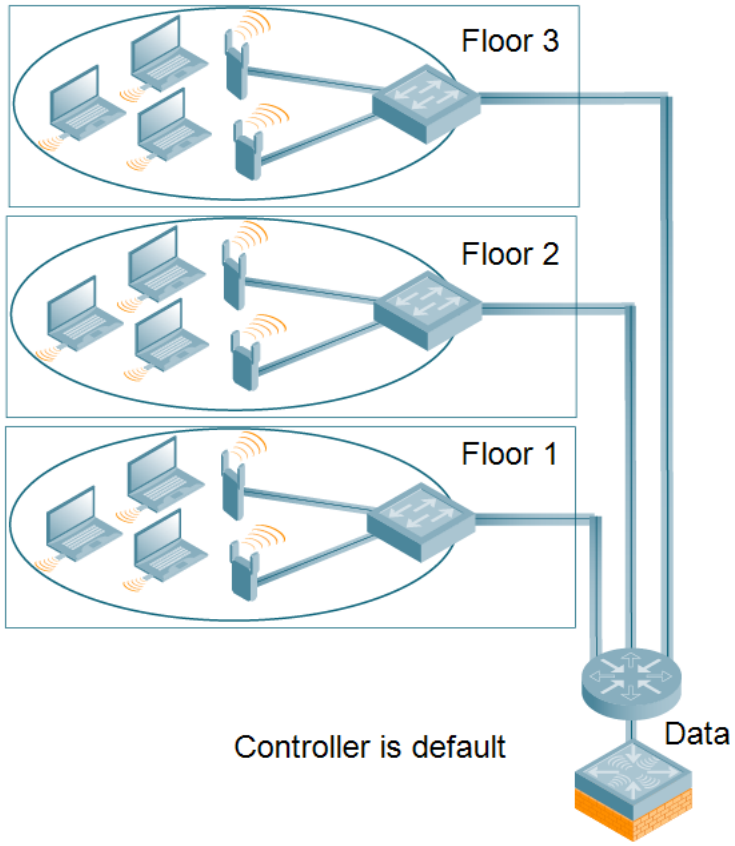
For this scenario, you must perform the following tasks:

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 managed device.

2. Connect the uplink port on the managed device to the switch or router interface. By default, all ports on the managed device are access ports and will carry traffic for a single VLAN.
3. Deploy APs. The APs will use the ADP to locate the managed device.
4. Configure the SSID(s) with VLAN 1 as the assigned VLAN for all users.

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

Figure 8 *APs All on One Subnet Different from Managed Device Subnets*



In this deployment scenario, the APs and the managed device are on different sub-networks and the APs are on multiple sub-networks. The managed device acts as a router for the wireless sub-networks (the managed device is the default gateway for the wireless clients). The uplink port on the managed device 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:

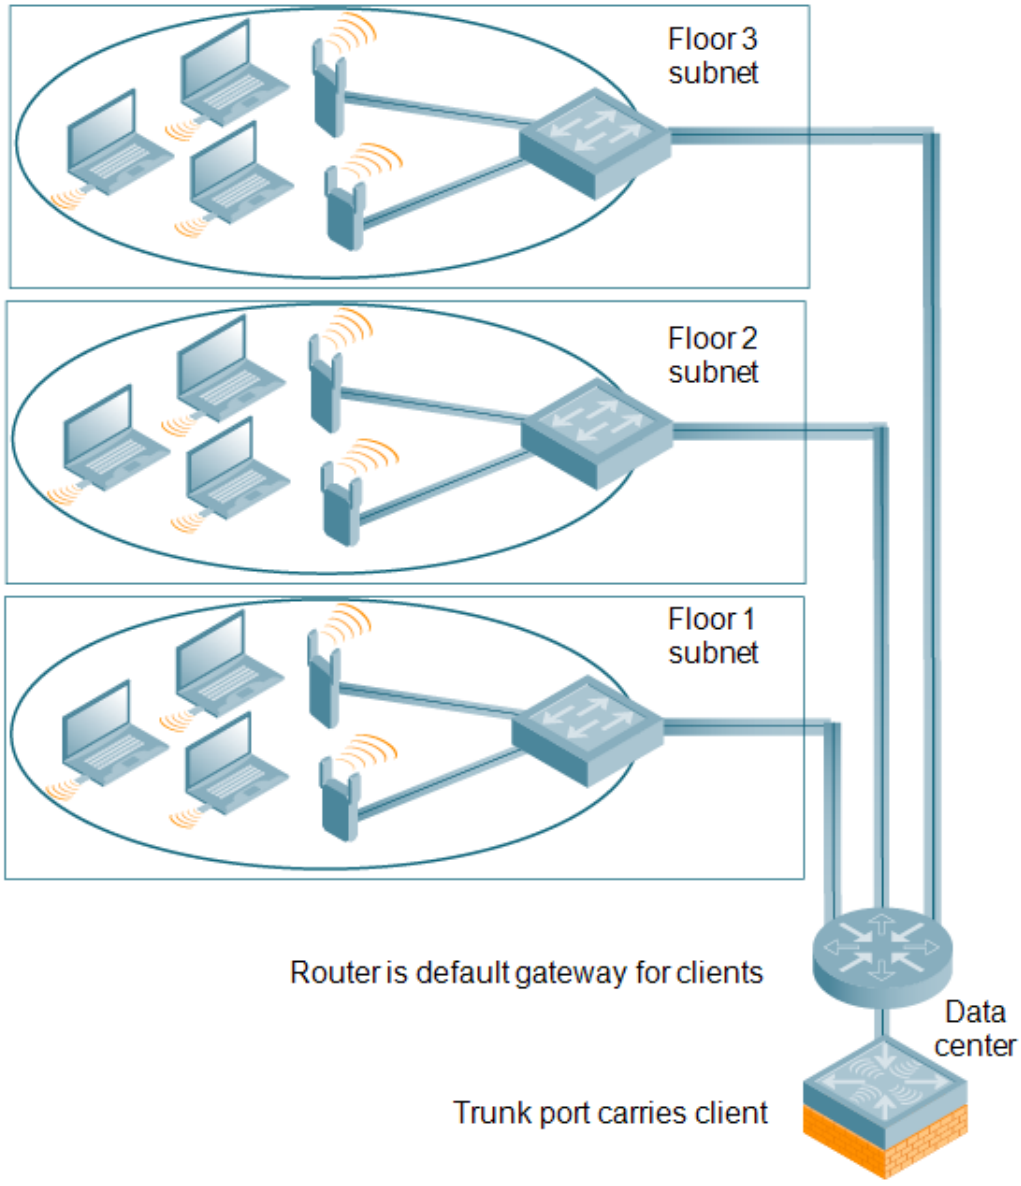
1. Run the initial setup wizard.
   - Set the IP address for VLAN 1.
   - Set the default gateway to the IP address of the interface of the upstream router to which you will connect the managed device.
2. Connect the uplink port on the managed device to the switch or router interface.
3. Deploy APs. The APs will use DNS or DHCP to locate the managed device.
4. Configure VLANs for the wireless sub-networks on the managed device.
5. Configure SSIDs with the VLANs assigned for each wireless sub-network.

NOTE

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

Deployment Scenario #3: APs on Multiple Different Subnets from Managed Devices

Figure 9 *APs on Multiple Different Subnets from Managed Devices*



In this deployment scenario, the APs and the managed device are on different sub-networks and the APs are on multiple sub-networks. There are routers between the APs and the managed device. The managed device 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.



The 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 managed device. Add the uplink port on the managed device 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 managed device. This gateway is the IP address of the router to which you will connect the managed device.
5. Configure the loopback interface for the managed device.
6. Connect the uplink port on the managed device to the switch or router interface.
7. Deploy APs. The APs will use DNS or DHCP to locate the managed device.
8. Now configure VLANs on the managed device for the wireless client sub-networks and configure SSIDs with the VLANs assigned for each wireless sub-network.

Managed Devices Configuration Workflow

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

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

The following workflow lists the tasks to configure a managed device. Click any of the links below for details on the configuration procedures for that task.

1. Connect the Managed Device to the Network.
2. Setting System Clock.
3. View current licenses and install new licenses.
4. For topologies similar to Deployment Scenario #3: APs on Multiple Different Subnets from Managed Devices on page 48), see Configuring VLANs to connect the managed device to your network. You do *not* need to perform this step if you are using VLAN 1 to connect the managed device to the wired network.
5. Configuring the Mobility Master IP Address. The managed device IP address is used by the managed device to communicate with external devices such as APs.
6. (Optional) Configuring the Loopback IP Address. You do *not* need to perform this step if you are using the VLAN 1 IP address as the managed device's IP address. Disable spanning tree on the managed device if necessary.

7. Configuring the Default Gateway for this managed device if you need to configure a trunk port between the managed device and another layer-2 switch (shown in Deployment Scenario #3: APs on Multiple Different Subnets from Managed Devices on page 48).
8. Trusted and Untrusted Ports and VLANs for this managed device.

Connect the Managed Device to the Network

To connect the managed device to the wired network, run the initial setup to configure administrative information for the managed device.

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 8.7.0.0 Quick Start Guide* and are referred to throughout this *chapter* as "initial setup."

This section describes the steps in detail.

Running the Initial Setup

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



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

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



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

The initial setup requires that you configure an IP address for the VLAN 1 interface, which you can use to access and configure the managed device 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 managed device upon completion of the initial setup.



The full setup dialog now provides flexibility to configure only IPv4 or IPv6 address, or a combination of both. If IPv6 address is used to terminate IPsec tunnel, then it is no longer mandatory to configure IPv4 address in master IP configuration in the setup dialog.

Connecting to the Managed Device after Initial Setup

After you complete the initial setup, the managed device reboots using the new configuration. (Refer to the *ArubaOS 8.7.0.0 Quick Start Guide* for information about using the initial setup.) You can then connect to and configure the managed device 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 managed device to enter the CLI. (See Management Access on page 938 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 managed device. You can then use one of the following access methods:
  - Use the VLAN 1 IP address to start an SSH session where you can enter CLI commands.
  - Enter the VLAN 1 IP address in a browser window to start the WebUI.
  - WebUI Wizards.



This chapter and the user guide in general focus on CLI and standard WebUI configuration examples. However, basic managed device configuration and WLAN or LAN creation can be completed using the alternative tasks (wizards) from within the WebUI. If you wish to use a configuration task, in the **Managed Network** node hierarchy, navigate to **Configuration > Tasks**, click the desired task, and follow the imbedded help instructions within the task.

7200 Series Controllers Port Behavior

The first two ports on the 7200 Series controllers, 0/0/0 and 0/0/1, are dual media ports and can be used for any purpose. Ports 0/0/2 through 0/0/5 are fiber-based ports that can be used for any purpose. If the fiber-based ports are connected with RJ45 or SFP transceivers, these ports can function as 1 Gbps ports. To access the controller, you can use port 0/0/0 to 0/0/5 when 0/0/2 through 0/0/5 are connected with RJ45 or SFP transceivers.

The following table describes the connector and speed supported for each physical interface of the 7200 Series controllers:

Table 10: *7200 Series Controllers Ports*

Port Type	Ports	Connector Type	Speed
10/100/1000 BASE-T Dual Media Ports	0/0/0–0/0/1	RJ45 or SFP	1 Gbps
10G BASE-X	0/0/2–0/0/5	SFP+	10 Gbps
		RJ45 or SFP	1 Gbps

Default Slot for USB Device

In 7000 Series and 7205 controllers, TRACES folder will not be automatically created when a USB device is connected to Slot 1. Ensure the USB device in Slot 0 as this is the default port where the TRACES folder is created.

Using the LCD Screen

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

Using the LCD Panel Mode

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

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

The LCD has four modes:

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

Table 11: *LCD Panel Mode: Boot*

Function or Menu Options	Display Output
Displays boot status	"Booting ArubaOS..."

Table 12: *LCD Panel Mode: LED Mode*

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

Table 13: *LCD Panel Mode: Status*

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

Table 14: *LCD Panel Mode: Maintenance*

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

Using the LCD and USB Drive

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



For more information on copying and transferring ArubaOS image and configuration files, see Managing Files on Managed Device on page 1004

Upgrading an Image

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

Uploading a Saved Configuration

1. Make a copy of a managed device configuration (with the .cfg file extension), and save the copied file with the name **Aruba_usb.cfg**.
2. Move the saved configuration file onto your USB drive into a directory named **/Arubaimage**.
3. Insert your USB drive into the managed device's USB slot. Wait for 30 seconds for the managed device to mount the USB.

4. Navigate to **Upload Config** in the **Maintenance** menu of the LCD. Confirm the upload (Y/N) and then wait for the upload to complete.
5. Execute a system reboot either from the LCD menu or from the command line to reload from the uploaded configuration.

For detailed upgrade and instruction, refer to the Upgrade chapter in the *ArubaOS 8.7.0.0 Release Notes*.

Disabling LCD Menu Functions

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

```
(host)[md](config) #lcd-menu
(host)[md](lcd-menu) #disable menu
```

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

```
(host)[md](lcd-menu) #disable menu maintenance ?
factory-default         Disable factory defaulting via LCD
halt-system             Disable system halt from LCD
media-eject             Disable media eject via LCD
reload-system           Disable system reload from LCD
upgrade-image           Disable image upgrade via LCD
upload-config           Disable config upload via LCD
```

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

```
(host) [md] #show lcd-menu
```

Configuring a VLAN to Connect to the Network

You must follow the instructions in this section only if you need to configure a trunk port between the managed device and another Layer-2 switch (shown in Deployment Scenario #3: APs on Multiple Different Subnets from Managed Devices on page 48).

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 managed device 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-managed device networks from a single location. For example, policies and virtual application configurations map users to different VLANs which may exist at a different managed device. This creates redundancy where one managed device has to back up many other managed devices. With the VLAN pool feature you can control your configuration globally.



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

- Assign to the VLAN the ports that you will use to connect the managed device 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 managed device 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 managed device.

The following sections provides step-by-step instructions to configure a VLAN and connect to the network.

Creating, Updating, and Viewing VLANs and Associated IDs

You can create and update a single VLAN or bulk VLANs using the WebUI or the CLI. See Configuring VLANs on page 83.



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

To view VLAN IDs in the CLI.

```
(host) [mynode] #show vlan
```

Creating, Updating, and Deleting VLAN Pools



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

You can create, update, and delete a VLAN pool using the WebUI or the CLI. See Configuring VLANs on page 83.

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

```
(host)[mynode](config) #vlan-name <name>
(host)[mynode](config) #vlan mygroup <vlan-ids>
```

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

```
(host)[mynode] #show vlan mapping
```

Assigning and Configuring the Trunk Port

The following procedure describes how to configure a Gigabit Ethernet port:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces > Ports** tab.
2. In the **Ports** section, click the port that will connect the managed device to the network.
3. Select **Trunk** from the **Mode** drop-down list.
4. Select a VLAN from the **Native VLAN** drop-down list.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure a Gigabit Ethernet port:

```
(host)[mynode](config) #interface gigabitethernet <slot>/<module>/<port>
(host)[mynode](config-submode) #switchport mode trunk
(host)[mynode](config-submode) #switchport trunk native vlan <id>
```

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

```
(host)[mynode] #show vlan
```

Configuring the Default Gateway

The following procedure describes how to configure the default gateway:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces > IP Routes** tab.
2. Click the **Static Default Gateway** accordion menu.
3. To add a new static gateway, click the **+** button below the static IP address list.
   a. Select **IPv4** or **IPv6** from the **IP version** drop-down list.
   b. In the **IP Address** field, enter an IP address with dot separators.
   c. In the **Cost** field, enter a value for the path cost.
   d. Click **Submit**.
4. You can define a dynamic gateway with the DHCP, PPPOE, or Cellular option by clicking the **Dynamic Default Gateway** accordion menu.
   a. In the **Dynamic Default Gateway** section, select the **DHCP**, **PPPoE** or **Cellular** check box to enable the corresponding dynamic gateway type. If you selected more than one dynamic gateway type, you must also define the cost for each gateway route. The managed device will first attempt to obtain a gateway IP address using the option with the lowest cost. If the managed device 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.
   b. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures the default gateway:

```
(host)[mynode](config) #ip default-gateway <ipaddr>|{import cell|dhcp|pppoe}|{ipsec
<name>} <cost> | mgmt | <nexthop>
```

Configuring the Loopback IP Address for the Managed Device

You must configure a loopback address if you are not using a VLAN ID address to connect the managed device to the network (see Deployment Scenario #3: APs on Multiple Different Subnets from Managed Devices on page 48).



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

If configured, the loopback address is used as the managed device's IP address. If you do not configure a loopback address for the managed device, the IP address assigned to the first configured VLAN interface IP address is considered. Generally, VLAN 1 is configured first and is used as the managed device's IP address. ArubaOS allows the loopback address to be part of the IP address space assigned to a VLAN interface. For example, if VLAN 5 interface on the managed device was configured with the IP address 10.3.22.20/24, the loopback IP address can be configured as 10.3.22.220.



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

STP is disabled by default on the managed device. STP ensures a single active path between any two network nodes, thus avoiding bridge loops. Disable STP on the managed device if you are not employing STP in your network.

The following procedure describes how to configure a loopback IP address:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > General** tab.
2. Click the **Loopback Interface** accordion menu.
3. Enter the **IPv4 address** and/or the **IPv6 address** in the corresponding text boxes.
4. Click **Submit**.
5. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > More** tab.
6. Click and expand **Spanning Tree**.
7. Click the **Spanning tree** toggle switch to enable this setting. By default, spanning tree is disabled.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.



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

11. In the **Mobility Master > host** node hierarchy, navigate to the **Maintenance > Software Management > Reboot** tab.
12. Select the **Save Current Configuration Before Reboot** check box.
13. Click **Reboot**.

The following CLI commands configure a loopback IP address:

```
(host)[mynode](config) #interface loopback ip address <A.B.C.D>
(host)[mynode](config) #no spanning-tree
(host)[mynode](config) #write memory
(host)[mynode](config) #reload
```

The managed device returns the following messages:

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

Enter **y** to reboot the managed device or **n** to cancel.

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

To verify that the managed device 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 managed device, or configure the managed device to use a NTP server to synchronize its system clock with a central time source. For more information about setting the managed device's clock, see Setting System Clock on page 1010.

Configuring the License Management with ASP

Starting from ArubaOS 8.4.0.0, ArubaOS License automation feature is supported where the Mobility Master obtains the licenses from Aruba Support Portal (ASP) or License Management Server automatically. The users need not manually add the licenses on the Mobility Master.

For the Mobility Master to obtain licenses, the users have to enter the ASP credentials using Mobility Master WebUI or the CLI only once.



The user can also assign new licenses to the Mobility Master using the WebUI instead of through Aruba Support Portal.

On-boarding ASP Licenses

Before signing on to ASP from Mobility Master, user must on-board the account from ASP, asp.arubanetworks.com.

Configuring license Management with Aruba Support Portal

The following procedure describes how to enable the ASP options:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration** > **System** > **General** tab.
2. Expand the **Aruba Support Portal (ASP)** section.
3. Enable the **Connect to ASP** option.
4. Enter the **Username** and **Password** to sign into Aruba Support Portal.
5. Click **Sign In**.
6. Click **Submit**.
7. To view the ASP license keys allotted to the Mobility Master, navigate to **Configuration** > **Licensing** > **License Inventory**.

You can also enable the ASP option using the following steps:

1. In the **Mobility Master** node hierarchy, navigate to **Mobility Master** > **Configuration** > **Licensing.**
2. Select **Aruba Support Portal (ASP)** option for **License management**.
3. Enter the **Username** and **Password** to sign in to Aruba Support Portal.
4. Click **Sign In**.
5. To view the ASP license keys allotted to the Mobility Master, navigate to **Configuration** > **Licensing** > **License Inventory**.

The following command creates, enables, and views the ASP profile:

Creating default ASP Profile

```
(host) [mm] (config) #asp-profile (can be executed in mm node only)
(host) [mm] (Aruba Support Portal Profile) #asp-enable
(host) [mm] (Aruba Support Portal Profile) #asp-licensing-enable
```

Signing On to ASP

```
(host) [mm] (config) #asp signon username <username>
```

Verifying the ASP sign-on status

```
(host) [mm] #show asp status
(host) [mm] #show asp standby status
```

Checking the ASP account used to login

```
(host) [mm] #show asp account-info
```

Registering or Claiming a license purchase and verify available licenses

```
(host) [mm] #license asp register-order <confirmationnumber> <ordernumber>
(host) [mm] #show license asp unallocated-lic
```

Allocating licenses

```
(host) [mm] #license asp allocate-lic ap <ap-num>
```



Allocation can be done for all license types at once or one by one

Verifying the PEFV licenses installed in Controllers

```
(host) [mm] #show license md-pefv-lic
```

Checking the total number of licenses allocated using ASP and Manual Licensing

```
(host) [mm] #show license summary
```

The following sections describe how to synchronize, view, allocate, and claim licenses:

Synchronizing Licenses between ASP and Mobility Master

Every successful sign-on attempt and also every time the Mobility Master is rebooted, the licenses between Aruba Support portal and Mobility Master are synchronized seamlessly.



Mobility Master synchronizes licenses from Aruba Support portal every 24 hours.

The following procedure describes how to synchronize the licenses from ASP to Mobility Master:

1. In the **Mobility Master** node hierarchy, navigate to **Mobility Master** > **Configuration** > **Licensing.**
2. Select **Aruba Support Portal (ASP)**.
3. Click **License Inventory** tab.
4. Click **Update now** to synchronize the activated licenses from ASP to Mobility Master.

The following CLI commands configure:

```
(host) [mm] #license asp get-allocated-lic
(host) [mm] #license asp get-md-pefv-lic
```

Viewing, Allocating, and Claiming Licenses

The following procedure describes how to view, allocate, or claim the license inventory:

1. In the **Mobility Master** node hierarchy, navigate to **Mobility Master** > **Configuration** > **Licensing.**
2. Click **License Inventory**.
3. The **License Inventory** tab lists detailed information about all the licenses used. It provides the following information:

- Type of License - the different type of licenses like AP, PEFNG, WEBCC, and so on.
- Description - the description of each type of license.
- Status - the status of the each type of license. For example, active license, not licensed, never licensed, and so on.
- Expiration - the expiration type of each license type.
- Total Activated - the sum of licenses allocated using ASP licensing method and licenses installed using Manual method.
- Available - the licenses which are successfully claimed or registered to ASP account using an Order or Confirmation Number.

4. To claim or register licenses, click **Claim** and enter **Order #** and **Confirmation #** and click **Submit**. The order Number and confirmation number is received through an email from Aruba Sales team after a successful license purchase.
5. To allocate or activate licenses, click **Allocate** and enter the number of licenses count for the license types in **ALLOCATE** column and click **Submit**.

For more information on licenses installation, refer to the *Aruba Mobility Master Licensing Guide*.

Offline Licensing feature

When a Mobility Controller Virtual Appliance stand-alone controller fails in a remote deployment, the backup stand-alone is brought up by deploying the OVA file but for the backup stand-alone controller should work with the same capacity and features of the failed stand-alone controller, it requires the same licenses.

This feature is supported only for Mobility Controller Virtual Appliance configured as a stand-alone controller.

In a scenario where the remote deployment has lost internet access or connection to the base, the user cannot activate the new license required for the backup standalone controller. In such a case, the offline licensing feature is used to activate new license using a Master Token Key (MTK).

The Master Token Key is generated by the user through LMS and this MTK is then, sealed in an envelope and provided to the user on a need basis. The MTK supports installing and activating MC-VA-XX licensing type, AP, PEF, and RFP licenses.



Webcc and ACR licenses cannot be installed through MTK.

For more information, see *Aruba Mobility Master licensing guide*.

Connecting the Managed Device to the Network

Connect the ports on the managed device 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 *Aruba Virtual Appliance Installation Guide* for details on the managed device for port LED and cable descriptions.



In many deployment scenarios, an external firewall is situated between various Aruba devices. External Firewall Configuration on page 831 describes the network ports that must be configured on the external firewall to allow proper operation of the network.

To verify that the managed device is accessible on the network:

- If you are using VLAN 1 to connect the managed device to the network (Deployment Scenario #1: Managed Device and APs on Same Subnet on page 46 and Deployment Scenario #2: APs All on One Subnet Different from Managed Device Subnet), ping the VLAN 1 IP address from a workstation on the network.
- If you created and configured a new VLAN (Deployment Scenario #3: APs on Multiple Different Subnets from Managed Devices on page 48), ping the IP address of the new VLAN from a workstation on the network.

Configuring User-Centric Network

Configuring your managed device and AP is done through either the WebUI or the CLI.

- WebUI is accessible through a standard Web browser from a remote management console or workstation. The WebUI includes configuration tasks that walk you through easy-to-follow configuration steps. Each task has embedded online help. The tasks are:
  - Provision New APs—basic AP configurations including LAN, Remote, LAN Mesh, and Remote Mesh deployment scenarios.
  - Controller—applicable only the first time the managed device is brought UP; basic managed device configuration including system settings, Control Plane security, and cluster settings.
  - Create a New WLAN—creating and configuring new WLANs and LANs associated with the “default” ap-group. Includes campus-only and remote networking.



Clicking **Cancel** from the tasks (wizards) return you to where you launched the tasks from. Any configuration changes you entered are not saved.

- The CLI allows you to configure and manage managed device. The CLI is accessible from a local console connected to the serial port on the managed device or through a Telnet or SSH session from a remote management console or workstation.



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

Replacing a Controller

The procedure below describes the steps to replace an existing stand-alone controller and/or a redundant controller. Best practice is to replace the backup controller first, and replace the active controller only after the new backup controller is operational on the network. When you remove the active controller from the network to replace it, the new backup controller takes over the active controller role. When you add a second controller to the network, the second controller automatically assumes the role of a backup controller.

For information on the Mobility Controller Virtual Appliance, refer to the *Aruba Virtual Appliance Installation Guide*.

Replacing an Returned Merchandise Authorization (RMA) Device

If the controller being replaced was returned to Aruba as an RMA device, the license keys on the RMA controller cannot be directly transferred to a new device, and must be regenerated.

To generate a new license key for a controller that is returned as an RMA:

1. Access the My Networking Portal at http://hpe.com/networking/mynetworking/.
2. Log in to the My Networking Portal using the HPE Passport.
3. Click **View licenses** or **Transfer licenses to new platform**. All available licenses are displayed.
4. Select the **>>** icon at the right end of the record to verify the license details before transferring it.
5. Click **Transfer License** at the bottom of the page.
6. Select a controller from the **AOS Controller Type** drop-down list.
7. Enter the serial number of the mobility controller in the **Serial number** text box; or enter the passphrase of the Mobility Master in the **PassPhrase** text box.
8. Select the license to be transferred.
9. Click **Transfer** at the bottom of the page. A new license key is generated, which you can apply to the controller.

Procedure Overview

The procedure to replace a backup or active controller consists of the following tasks:

1. Step 1: (Optional) Change the VRRP Priorities for a Redundant Master Pair on page 62
2. Step 2: Back Up the Flash File System on page 62
3. Step 3: Stage the New Controller on page 63
4. Step 4: Add Licenses to the New Controller on page 63
5. Step 5: Backup Newly Installed Licenses on page 63
6. Step 6: Import and Restore the Flash Backup on page 63
7. Step 7: Restore Licenses on page 64
8. Step 8: Reboot the Controller on page 64
9. Step 9. (Optional) Modify the Host Name on page 64
10. Step 10: Save your Configuration on page 65
11. Step 11: Remove the Existing Controller on page 65



If your controller does not have any manually added licenses, skip steps 3, 4, and 6 of the following procedure.

Step 1: (Optional) Change the VRRP Priorities for a Redundant Master Pair

If your deployment uses VRRP to define the primary Mobility Master in a pair of redundant Mobility Masters and you are replacing only the primary Mobility Master, you must change the VRRP priority levels of the controllers so that the primary Mobility Master has a lower priority than the backup Mobility Master. This will allow the configuration from the backup Mobility Master to be copied to the new Mobility Master, and prevent an old or inaccurate configuration from being pushed to the managed devices.

Step 2: Back Up the Flash File System

To start the migration process, access the backup controller or the Mobility Master being replaced and create a backup of the flash file system. You can create a backup file using the WebUI or command-line interfaces.

To create a flash backup from the command-line interface, access the active controller and issue the **backup flash** command.

To back up the flash from the WebUI, log in to the current backup controller or active controller and create a flash backup using the procedure below:

1. In the **Mobility Master** node hierarchy, select the device and navigate to the **Maintenance > Configuration Management > Backup** tab.
2. For the **Select what to backup** option, select **Flash**.
3. Click **Create Backup**. A confirmation message (**Backup saved successfully**) is displayed.
4. Click **Copy Backup** to create a copy of the backup file. By default, the flash backup file is named **flashbackup.tar.gz**.
5. In the **Select source file** drop-down list, select **Flash file system**.
6. In the **File name** drop-down list, select the relevant file name.
7. In the **Select destination file** drop-down list, select one of the server options to move the flash backup off the controller.
8. In the **File name** text box, enter the name of the Flash backup file to be exported.
9. Click **Copy**. A confirmation message (**Files copied successfully**) is displayed.

Step 3: Stage the New Controller

The next step in the procedure is to stage the new backup controller or active controller with basic IP connectivity. Power up the new controller, connect a laptop computer to the controller's serial port, and follow the prompts to configure basic settings, such as the controller name, role, VLAN, gateway, country code, and time zone.

Step 4: Add Licenses to the New Controller

To replace a controller with manually added licenses, you will need to transfer those licenses to the new controller as part of the replacement process.

Use the **license add** command in the command-line interface. Alternatively, in the **Mobility Master** node hierarchy, navigate to the **Configuration > License** page to add new or transferred licenses to the new controller.



Do **not** reboot the controller at the end of this step. Do not save the configuration or write it to memory. Reboot only after the flash memory and the licenses have been restored.

Step 5: Backup Newly Installed Licenses

Use the **license export** command in the command-line interface to back up the newly installed licenses to the backup license database.

```
(host)[mynode] #license export <filename>
```



Do **not** reboot the controller at the end of this step. Do not save the configuration or write it to memory. Reboot only after the flash memory and the licenses have been restored.

Step 6: Import and Restore the Flash Backup

Import and restore the backup flash file system from the original controller to the new controller.



Do **not** reboot the controller at the end of this step. Do not save the configuration or write it to memory. Reboot only after the flash memory and the licenses have been restored.

The following procedure describes how to import and restore a flash backup:

1. Access the new controller.
2. In the **Mobility Master** node hierarchy, select the device and navigate to the **Diagnostics > Technical Support > Copy Files** tab.
3. In the **Select source file** drop-down list, select any of the source options.
4. In the **File name** drop-down list, select the relevant source file name.
5. In the **Select destination file** drop-down list, select **Flash file system**.
6. In the **File name** text box, enter the name of the Flash backup file.
7. Click **Copy**. By default, the flash backup file is named **flashbackup.tar.gz**. A confirmation message (**Files saved successfully**) is displayed.
8. Next, to restore the backup of the flash file system, navigate to **Mobility Master > host** node hierarchy, navigate to **Maintenance > Configuration Management > Restore** tab.
9. Select **Flash** and click **Restore**. A confirmation message (**Flash restored successfully**) is displayed.

To import and restore a flash backup file using the command-line interface, use the **copy** and **restore flash** commands. The following example copies a backup file from a USB drive:

```
(host)[mynode] #copy usb: Partition 1 flashbak2_3600.tar.gz flash: flashbackup.tar.gz
....File flashbak2_3600.tar.gz copied to flash successfully.
(host)[mynode] #restore flash
```

Step 7: Restore Licenses

Execute the **license import** command in the command-line interface to import licenses from the license database to the new controller.

```
(host)[mynode] #license import <filename>
```



Do **not** save the configuration or write to memory at the end of this step.

Step 8: Reboot the Controller

When all the licenses have been restored, issue the **reload** command in the command-line interface. Alternatively, in the **Mobility Master** node, navigate to **Maintenance > Software Management > Reboot** in the WebUI to reboot the new controller. After rebooting, the controller should not be on the network (or a reachable subnet) with the controller it will replace. This is to prevent a possible IP address conflict.



Do **not** save the configuration or write to memory at the end of this step.

```
(host)[mynode] #reload
Do you want to save the configuration(y/n): n
Do you really want to restart the system(y/n): y
System will now restart!
```

Step 9. (Optional) Modify the Host Name

Execute the **hostname** command in the command-line interface to give the new controller a unique host name. (The flash restoration process gives the new controller the same name as the existing controller.)



Do **not** save the configuration or write to memory at the end of this step.

Step 10: Save your Configuration

Now, you must save the configuration settings on the new controller. Execute the **write memory** command in the command-line interface, or in the Web UI navigate to the **Managed Network** node hierarchy, click the **Configuration** tab and select **Pending Changes** at the top of the WebUI page.

Step 11: Remove the Existing Controller

If you are only replacing a backup controller, remove the existing backup controller and then connect the replacement controller to the network. If you are replacing both an active controller and a backup controller, replace the backup controller first.

When the active controller is removed from the network, the backup controller immediately assumes the role of active controller, and all active APs associate to the new active controller within a few seconds. Therefore, when you add another controller to the network, it will, by default, assume the role of a backup controller.

If you changed the VRRP priorities of your redundant Mobility Master prior to replacing the primary Mobility Master, you may wish to change them back once the new primary Mobility Master is active on the network.



When the new controller uses the same IP address of the controller that is being replaced, it is recommended to issue the **apboot** command after the APs connect to the controller. If the APs were already rebooted before connecting to the controller, the **apboot** command need not be issued.

Chapter 5
Control Plane Security

ArubaOS supports secure IPsec communications between a managed device and campus APs or remote APs using public-key self-signed certificates created by each Mobility Master. The managed device certifies its APs by issuing them certificates.

If the Mobility Master has any associated managed device, the Mobility Master sends a certificate to each managed device, which in turn sends certificates to their own associated APs. If a managed device is unable to contact the Mobility Master to obtain it's own certificate, it will not be able to certify the APs, and those APs can not communicate with their managed device until Mobility Master-managed device communication has been re-established. You create an initial CPsec configuration when you first configure the managed device using the initial setup wizard. The ArubaOS initial setup wizard enables CPsec by default, so it is very important that the managed device be able to communicate with the Mobility Master when it is first provisioned.

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

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

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

Topics in this chapter include:

- Control Plane Security Overview on page 66
- Configuring Control Plane Security on page 67
- Managing AP Whitelists on page 68
- Whitelist DB Optimization on page 76
- Configuring Networks with a Backup Mobility Master on page 76
- Replacing a Controller on a Multi-Controller Network on page 77
- Troubleshooting Control Plane Security on page 78

Control Plane Security Overview

Controllers using CPsec send certificates to APs that you have identified as valid APs on the network. If you want closer control over each AP that is certified, you can manually add individual campus and remote APs to the secure network by adding each AP's information to the whitelists when you first run the initial setup wizard. If you are confident that all APs currently on your network are valid APs, then you can use the initial setup wizard to configure automatic certificate provisioning to send certificates

from the controller to each campus or remote AP, or to all campus and remote APs within specific ranges of IP addresses.

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

Configuring Control Plane Security

When you initially deploy the controller, you create your initial CPsec configuration using the initial setup wizard. These settings can be changed at any time using the WebUI or CLI. The following procedure describes how to create the initial CPsec configuration.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > CPSec** tab.
2. Select the **Control Plane Security** accordion.
3. Click the **Enable CPSEC** toggle switch to enable this setting.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

To enable auto cert provisioning:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > CPSEC** tab.
2. Select the **Control Plane Security** accordion.
3. Click the **Enable CPSEC** toggle switch to enable this setting.
4. Click the **Enable auto cert provisioning** toggle switch to allow AP's from specified ranges.
5. Click the **Only accept APs from specified ranges** toggle switch to enable this setting.
    a. Click **+** in **Address ranges for Auto Cert Provisioning** table. The **New Address Range** window is displayed.
    b. Enter the IPv4 or IPv6 address in the **Start address IPv4 or v6** and **End address IPv4 or v6** fields.
    c. Click **OK**.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes.**

The Mobility Master generates its self-signed certificate and begins distributing certificates to campus APs and any managed devices 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.

Table 15: *Control Plane Security Parameter*

Parameter	Description
Enable CPSEC	Select **enable** or **disable** to turn the control plane security feature on or off. This feature is enabled by default.
Enable auto cert provisioning	When you enable the control plane security feature, you can toggle this switch to turn on automatic certificate provisioning. When you enable this feature, the controller attempts to send certificates to all associated campus APs. Auto certificate provisioning is disabled by default. **NOTE:** If you do not want to enable automatic certificate provisioning the first time you enable control plane security on the controller, you must identify the valid APs on your network by adding those to the campus AP whitelist. For details, see Managing AP Whitelists on page 68. After you have enabled automatic certificate provisioning, you must select **Only accept APs from specified ranges.**
Only accept APs from specified ranges	Enabling this option will let you automatically certify APs within a select range of IP addresses.
Address ranges for Auto Cert Provisioning	The **Address ranges for Auto Cert Provisioning** section allows you to send certificates to a group of campus or remote APs within a range of IP addresses. Click + to specify the start and end IP address of the range. Repeat this procedure to add additional IP ranges to the list of allowed addresses. If you enable both control plane security and auto certificate provisioning, all APs in the address list receives automatic certificate provisioning. Remove a range of IP addresses from the list of allowed addresses by selecting the IP address range from the list and clicking **Delete.**

The following commands configure CPsec on a managed device or Mobility Master:

```
(host) [md] (config) #control-plane-security
(host) [md] (Control Plane Security Profile) #auto-cert-allow-all
(host) [md] (Control Plane Security Profile) #auto-cert-allowed-addrs <start> <end>
(host) [md] (Control Plane Security Profile) #auto-cert-prov
(host) [md] (Control Plane Security Profile) #cpsec-enable
The following command enables you to view the current CPsec settings:
(host) [md] (config)#show control-plane-security
```

Managing AP Whitelists

Campus or Remote APs appear as valid APs in the Campus AP or Remote AP whitelists when you manually enter their information into the Campus AP or Remote AP whitelists using the WebUI or CLI of a controller. Also, the Campus APs or Remote APs appear as valid APs after a controller sends a certificate to an AP as part of automatic certificate provisioning and the AP connects to the controller over a secure tunnel. APs that are not approved or certified on the network are included in the Campus AP whitelists, but these APs appear in an unapproved state.

Use the AP whitelists to grant valid APs secure access to the network or to revoke access from suspected rogue APs. When you revoke or remove an AP from the Campus AP or Remote AP whitelists on a controller that uses CPsec, that AP will not able to communicate with the controller again, unless the AP obtains a new certificate.

The following sections discuss the procedures to manage AP whitelists:

Adding an AP to the Campus or Remote AP Whitelists

You can add an AP to the Campus AP or Remote AP whitelists using the WebUI or CLI. The following procedure describes the steps to add an AP to the Campus AP or Remote AP whitelist:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Access Points > Whitelist** tab.
2. Click **Campus AP Whitelist** or **Remote AP Whitelist** tab.
3. Click **+**.
4. Define the following parameters for each AP you want to add to the AP whitelist:

Table 16: *AP Whitelist Parameters*

Parameter	Description
Campus AP whitelist configuration parameters	
MAC address	MAC address of campus AP that supports secure communications to and from its controller.
AP name	Name of the campus AP. If you do not specify a name, the AP uses its MAC address as AP name.
AP group	Name of the AP group to which the campus AP is assigned. If you do not specify an AP group, the AP uses default as its AP group.
Description	Brief description of the campus AP.
Remote AP whitelist configuration parameters	
MAC address	MAC address of the remote AP, in colon-separated octets.
AP name	Name of the Remote AP. If you do not specify a name, the AP uses its MAC address as AP name.
AP group	Name of the AP group to which the Remote AP is assigned.
Description	Brief description of the Remote AP.
IPv4 address	IPv4 address of the Remote AP.
IPv6 address	IPv6 address of the Remote AP.

5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command adds an AP to the Campus AP whitelist:

```
(host) [mynode] (config) #whitelist-db cpsec add mac-address <address>
ap-group <ap_group>
ap-name <ap_name>
description <description>
```

The following CLI command adds an AP to the Remote AP whitelist:

```
(host) [mynode] (config) #whitelist-db rap add mac-address <mac-address>
ap-group <ap-group>
 ap-name <ap-name>
 description <description>
 full-name <name>
 remote-ip <inner-ip-adr>
 remote-ipv6 <ipv6 address>
```

Viewing AP Whitelist Entries

The WebUI displays the table of entries in the selected AP whitelist. The table of entries page displays a list of AP whitelist entries.

The **Configuration > Access Points > Whitelist** tab displays the list of the Campus AP whitelists by default. To view the list of Remote AP whitelists, click **Remote AP whitelist**.

The Remote AP whitelist entries page displays only the information you can manually configure. The Campus AP whitelist entries page displays both user-defined settings and additional information that are updated when the status of a Campus AP changes.

Table 17: *Campus AP Parameters*

Parameter	Description
Status	Displays the status of the AP whitelist entry.
Revoke text	Brief description for revoking the campus AP.
Approved	Approval status of the campus AP.
Updated	Time and date of the last AP status update.

To view information about the Campus AP and Remote AP whitelists using the CLI, issue the following commands:

```
(host) [mynode] #show whitelist-db cpsec
Control-Plane Security Whitelist-entry Details
----------------------------------------------
MAC-Address        AP-Group  AP-Name  Enable   State                    Cert-Type
Description  Revoke Text  Last Updated
-----------        --------  -------  ------   -----                    ---------     -
----------  -----------  ------------
6c:f3:7f:cc:42:25                     Enabled  certified-factory-cert  factory-cert
                       Thu Jul  7 03:42:21 2016
9c:1c:12:c0:7c:a6  default   san225   Enabled  certified-factory-cert  factory-cert
                       Wed Aug  3 10:34:13 2016
24:de:c6:ca:94:ba                     Enabled  certified-factory-cert  factory-cert
                       Fri Apr 22 06:28:46 2016
94:b4:0f:c0:cc:42                     Enabled  certified-factory-cert  factory-cert
                       Fri Aug  5 06:54:43 2016
18:64:72:cf:e6:9c                     Enabled  certified-factory-cert  factory-cert
                       Tue Aug  9 07:35:41 2016
ac:a3:1e:c0:e6:82                     Enabled  certified-factory-cert  factory-cert
                       Wed Aug 10 09:12:23 2016
ac:a3:1e:cd:36:84                     Enabled  certified-factory-cert  factory-cert
                       Fri Jun 17 05:50:02 2016
ac:a3:1e:c0:e6:9a                     Enabled  certified-factory-cert  factory-cert
                       Thu May 26 06:31:13 2016
Total Entries: 8

(host) [mynode] #show whitelist-db cpsec-status
```

```
My Mac-Address                00:1a:1e:00:1a:b8
My IP-Address                 10.15.28.16
Master IP-Address             10.15.28.16
Switch-Role                   Master
Whitelist-sync is disabled
Entries in Whitelist database
Total entries:                5
Approved entries:             0
Unapproved entries:           2
Certified entries:            2
Certified hold entries:       1
Revoked entries:              0
Marked for deletion entries:  0
Current Sequence Number:    147

(host) [mynode] #show whitelist-db rap
Entries in Whitelist database
Total entries:                0
Revoked entries:              0
Marked for deletion entries:  0
AP Entries: 4
```

Modifying an AP in the Campus AP Whitelist

Use the following procedures to modify the AP group, AP name, certificate type, state, description, and revoked status of an AP in the Campus AP whitelist using the WebUI or CLI.

The following procedure describes the steps to modify an AP in the Campus AP whitelist:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Access Points > Whitelist** tab.
2. Click **Campus AP Whitelist** tab.
3. Select the check box of the AP that you want to modify.
4. Modify the settings of the selected AP. Some of the following parameters are available when adding an AP to the Campus AP whitelist.
   - **AP name**: The name of the Campus AP. If you not specify a name, the AP uses its MAC address as a name.
   - **AP group**: The name of the AP group to which the Campus AP is assigned.
   - **Description**: Brief description of the Campus AP.
   - **Status**: Select **Revoked** or **Accepted**.
   - **Revoked string**: Enter a value for this string.
5. Click **Submit** to update the Campus AP whitelist entry with its new settings.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands modify an AP in the Campus AP whitelist:

```
(host) #whitelist-db cpsec modify mac-address <name>
        ap-group <ap_group>
        ap-name <ap_name>
        cert-type {switch-cert|factory-cert}
        description <description>
        mode {disable|enable}
        revoke-text <revoke-text>
        state {approved-ready-for-cert|certified-factory-cert}
```

Revoking an AP from the Campus AP Whitelist

You can revoke an invalid or rogue AP either by modifying its revoke status (as described in Modifying an AP in the Campus AP Whitelist on page 71) or by directly revoking it from the Campus AP whitelist without modifying any other parameter. When revoking an invalid or rogue AP, enter a brief description why the AP is being revoked. When you revoke an AP from the Campus AP whitelist, the Campus AP whitelist retains the information of the AP. To revoke an invalid or rogue AP and permanently remove it from the whitelist, delete that entry.

You can revoke an AP from the Campus AP whitelist using the WebUI or CLI.

The following procedure describes the steps to revoke an AP from the Campus AP whitelist:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Access Points > Whitelist** tab.
2. Click **Campus AP Whitelist** tab.
3. Click on the check box next to the AP you want to revoke and click **Revoke**. The **Revoke** window is displayed.
4. Enter a brief description of why the AP is being revoked in the **Revoke text** field.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command revokes an AP via the Campus AP whitelist:

```
(host) [mynode] (config) #whitelist-db cpsec revoke mac-address <name> revoke-text
<comment>
```

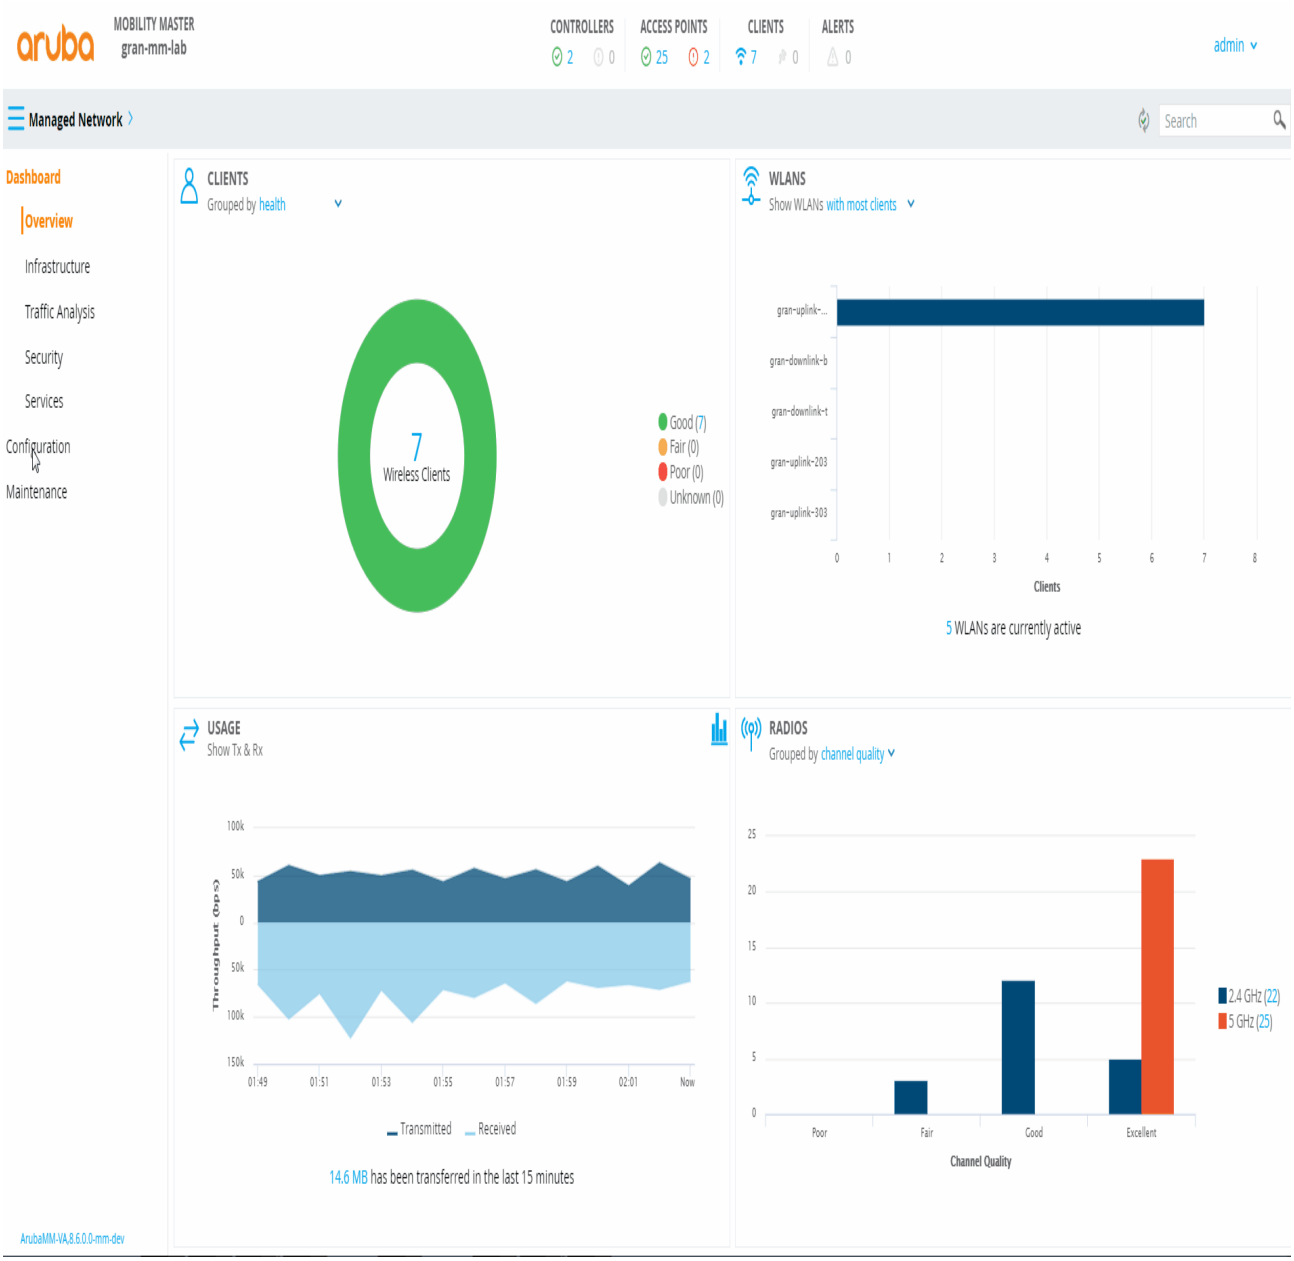
Deleting an AP from the Campus AP Whitelist

Before deleting an AP from the Campus AP whitelist, verify that auto certificate provisioning is either enabled or disabled only for IP addresses that do not include the AP being deleted. If you enable automatic certificate provisioning for an AP that is still connected to the network, you cannot delete it from the Campus AP whitelist; the controller immediately re-certifies the AP and re-creates its whitelist entry.

You can delete an AP from the Campus AP whitelist using the WebUI or CLI. The following procedure describes the steps to delete an AP from the Campus AP whitelist:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Access Points > Whitelist** tab.
2. Click **Campus AP Whitelist** tab.
3. Select the check box of the AP that you want to delete, then click **Delete**.
4. Click **Delete**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following animation displays how to delete an AP from the Campus AP whitelist in the WebUI:



The following CLI command deletes an AP from the Campus AP whitelist:

```
(host) [mynode] (config) #whitelist-db cpsec del mac-address <name>
```

Purging a Campus AP Whitelist

Before adding a new managed device to a network using CPsec, purge the campus AP whitelist on the new managed device. To purge a Campus AP whitelist, issue the following command:

```
(host) [mynode] (config) #whitelist-db cpsec purge
```

Offloading a Controller Whitelist to ClearPass Policy Manager

This feature allows to externally maintain AP whitelist in a ClearPass Policy Manager server. The controller, if configured to use an external server, can send a RADIUS access request to a ClearPass Policy Manager server. The MAC address of the AP is used as a username and password to construct the

access request packet. The ClearPass Policy Manager server validates the RADIUS message and returns the relevant parameters for the authorized APs.

The following supported parameters are associated with the following Vendor Specific Attributes (VSAs). The ClearPass Policy Manager server sends them in the RADIUS access accept packet for authorized APs:

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

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

- ap-group: The default ap-group is assigned to the AP.
- ap-name: The MAC address of the AP is used as the AP name.

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

ArubaOS now provides support for ClearPass Policy Manager to whitelist Remote APs in a cluster environment. You can configure ClearPass Policy Manager as an external server that authenticates Remote APs using the MAC address of Remote APs. The Remote APs are authenticated by maintaining whitelist entries in ClearPass Policy Manager, and the cluster inner IP addresses are assigned on the Mobility Master. Hence, the inner IP address assignment is centralized and forwarded to the associated managed devices in the cluster.

The following procedure describes the steps to assign a ClearPass Policy Manager server to a Remote AP:

1. Configure a ClearPass Policy Manager server using the WebUI:

a. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** tab.
b. Click **+** in the **Server Groups** table.
c. In the **Add Server Group** window, enter the server group name in the **Name** field.
d. Click **Submit**.
e. Select the server group created.
f. Click **+** in the **Server Group > <name>** table.
g. To assign an existing server as the ClearPass Policy Manager server,
   - Select **Add existing server** option.
   - Choose a server from the list.
   - Click **Submit**.
h. To create a new ClearPass Policy Manager server,
   - Select **Add new server** option.
   - Enter/ Select appropriate values in the following fields:
   - **Name,**
   - **IP address / hostname**
   - **Type**
   - Click **Submit**.
   - Select the new server created in the **All Servers** table.

- Under **Server Options**, enter a value in the **Shared Key** field and re-enter the value in the **Retype key** field.
- Click **Submit**.

i. Click **Pending Changes**.
j. In the **Pending Changes** window, select the checkbox and click **Deploy Changes**.

2. In the **Mobility Master** node hierarchy, navigate to the **Configuration** > **System** > **Profiles** tab.
3. In the **All profiles** list, select **Wireless LAN > VPN Authentication> default-rap> Server Group.**
4. Select the ClearPass Policy Manager server from the **Server Group** drop-down list.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

To assign a ClearPass Policy Manager server to a Remote AP that was initially an Instant AP:

1. Ensure that a ClearPass Policy Manager server is configured on the controller.
2. In the **Mobility Master** node hierarchy, navigate to the **Configuration** > **System** > **Profiles** tab.
3. In the **All profiles** list, select **Wireless LAN > VPN Authentication> default-iap> Server Group.**
4. Select the ClearPass Policy Manager server from the **Server Group** drop-down list.
5. Click **Save**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following commands add a ClearPass Policy Manager server to a Remote AP:

Configure a RADIUS server with ClearPass Policy Manager server as host address. In this example **cppm-rad** is the ClearPass Policy Manager server name and **cppm-sg** is the server group name.

```
(host) [md] (config) #aaa authentication-server radius cppm-rad
(host) [md] (RADIUS Server "test") # host 1.1.1.1
```

Run the following commands to add this server to a server group:

```
(host) [md] (config) #aaa server-group cppm-sg
(host) (Server Group "cppm-sg") #auth-server cppm-rad
```

Run the following commands to add this server group to the **default-rap** vpn profile:

```
(host) [md] (config) #aaa authentication vpn default-rap
(host)(VPN Authentication Profile "default-rap") #server-group cppm-sg
```

Run the following command to configure the Remote AP inner IP pool on the Mobility Master for cluster deployment :

```
(host) [mynode] (config) #lc-rap-pool rap-cluster 3.1.1.3 3.1.1.10
```

Important Points to Remember

- The **lc-rap-pool** command currently supports only IPv4 address in a cluster environment.
- In the cluster environment, the managed device does not use the IP address received from ClearPass Policy Manager, and tries to obtain the cluster inner IP address from Remote AP inner IP pool for cluster deployment **(lc-rap-pool**) configured on the Mobility Master. If the managed device fails to obtain the inner IP address, the Remote AP does not establish IKE/IPsec tunnel with the managed

device. The whitelist entries are automatically generated after successful authentication and IP assignment from the Remote AP inner IP pool.

- When the Remote AP goes down on all cluster members, both the managed device and Mobility Master delete the Remote AP whitelist entries that are generated automatically.

Whitelist DB Optimization

In addition to the existing push-based model that syncs whitelist entries to managed devices when they are updated, deleted, or revoked from Mobility Master. The Mobility Master introduces a pull-based sync mechanism for the whitelist database, in which AP whitelist entries are only synced to the managed devices that require the entry. The pull-based sync mechanism is used when a Remote AP or CPsec AP terminates on a managed device or if a network is down during a whitelist push, which can prevent messages from going through to the managed devices. The managed device can use this as a fallback mechanism to periodically check if it is in sync with the Mobility Master. If a mismatch is detected, the managed device pulls the new entry from Mobility Master. All whitelist entries are configured from a centralized location on the Mobility Master and synced to appropriate managed devices. Entries can also be configured directly on a managed device for debugging purposes. However, these changes are not synced back to the Mobility Master or any other managed device.

This whitelist-DB optimization provides the following enhancements on Mobility Master:

- Reduced memory footprint.
- Increased performance on the Mobility Master and managed devices.
- Scalability and support for over 1000 managed devices and 10,000 APs on a Mobility Master.
- Scalability and support for managed devices with varying AP capacities.
- Simplified debugging process, as corrupt entries are no longer synced to every managed device on a given Mobility Master.



Changes made to the whitelist-DB can only be applied to the postgres database and are not backwards-compatible.

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

```
(host) #show whitelist-db seq-pendlist
```



In a Mobility Master, only a global list of whitelist entries are available. To view the entries specific to a managed device, login into the particular device to view the whitelist specific to the device.

Configuring Networks with a Backup Mobility Master

This section describes the configuration with a backup Mobility Master.

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

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

- **Manual Synchronization**: Issue the **database synchronize** command to manually synchronize databases from your primary Mobility Master to the backup Mobility Master.
- **Automatic Synchronization**: Schedule automatic database backups using the **database synchronize period** command in configuration mode.



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

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 Mobility Master or a cluster of Mobility Masters, and whether or not the controller has a backup.

Replacing Controllers in a Single Mobility Master Network

Use the procedures in this section to replace a Mobility Master or managed device in a network environment with a single Mobility Master.

Replacing a Managed Device

Follow the steps below to replace a managed device in a single Mobility Master network:

1. Disconnect the managed device from the network.
2. If you plan on moving the managed device to another location on the network, purge the campus AP whitelist on the managed device.
3. Access the CLI on the old managed device and issue the **whitelist-db cpsec purge** command.
4. Install the new managed device, but do not connect it to the network. If the managed device has been previously installed on the network, you must ensure that the new managed device has a clean whitelist.
5. Purge the managed device whitelist by executing the **whitelist-db cpsec purge** command on the new managed device.
6. Once the managed device has a valid CPsec certificate and configuration, the managed device receives the campus AP whitelist from the Mobility Master and starts certifying approved APs.
7. APs associated with the new managed device reboots and creates new IPsec tunnels to the controller using the new certificate keys.

Replacing a Redundant Mobility Master

The CPsec feature requires you to synchronize databases from the primary Mobility Master to the backup Mobility Master at least once after the network is up and running. This ensures that all certificates, keys, and whitelist entries are synchronized to the backup Mobility Master. Because the AP whitelist may change periodically, you should regularly synchronize these settings to the backup Mobility Master. For details, see Configuring Networks with a Backup Mobility Master on page 76.

When you install a new backup Mobility Master, *you must add it as a lower priority* controller than the existing primary Mobility Master. After you install the backup Mobility Master on the network, synchronize the database from the existing primary Mobility Master to the new backup Mobility Master

to ensure that all certificates, keys, and whitelist entries required for CPsec are added to the new backup Mobility Master configuration. If you want the new Mobility Master to act as the primary Mobility Master, you can increase that Mobility Master's priority *after* the settings have been synchronized.



The CPsec settings of a controller does not change if you upgrade the controller running ArubaOS 6.x to ArubaOS 8.0.0.0. If CPsec was already enabled, then it remains enabled after the upgrade, however if CPsec was not enabled previously and you want to use this feature after upgrading, then you must manually enable CPsec.

Troubleshooting Control Plane Security

Follow the procedures below to identify and troubleshoot CPsec issues:

Identifying Certificate Problems

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

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

Verifying Certificates

If you are unable to configure the CPsec security feature, verify that its TPM and factory-installed certificates are present and valid by accessing the controller's CLI and issuing the **show tpm cert-info** command. If the controller has a valid certificate, the output of the command appears similar to the output in the example below.



This command works only on hardware controllers.

```
(host) #show tpm cert-info
====================================
TPM manufacturing factory certificate
====================================
subject= /CN=BA0003137::00:1a:1e:00:89:b8
issuer= /DC=com/DC=arubanetworks/DC=ca/CN=DEVICE-CA1
serial=2E1DF0D10000004C8EE7
notBefore=Aug  6 22:50:04 2013 GMT
notAfter=Sep 14 03:21:14 2032 GMT
====================================
Generated Factory certificate
====================================
subject= /CN=BA0003137::00:1a:1e:00:89:b8/L=SW
```

```
issuer= /CN=BA0003137::00:1a:1e:00:89:b8
serial=2E1DF0D10000004C8EE7
notBefore=Aug  6 22:50:04 2013 GMT
notAfter=Sep 14 03:21:14 2032 GMT
```

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

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

Disabling Control Plane Security

If you disable CPsec on a Mobility Master or managed device, all APs connected to that controller reboot then reconnect to the controller over a clear channel.

If you disable CPsec for a managed device, APs directly connected to the managed device reboot and reconnect to the managed device over a clear channel.

Verifying Whitelist Synchronization

To verify if the campus AP whitelist is downloaded from the Mobility Master to managed devices, check the sequence numbers on the Mobility Master and managed device whitelists.

The sequence number value on a Mobility Master should be the same as the sequence number on the managed device.

Rogue APs

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

This section gives an overview of ArubaOS WLAN. It describes the procedures to configure a basic WLAN and define VLANs and ports. It also describes how to enable advanced WLAN, optional WLAN, and VLAN optimization features.

Click any of the following links to configure a campus WLAN:

- Campus WLAN Workflow on page 80
- Configuring VLANs on page 83
- Trusted and Untrusted Ports and VLANs on page 91

Click the following links to configure a basic network:

- Assign an IP Address to a VLAN on page 92
- Configuring Trusted or Untrusted Ports and VLANs on page 96
- Configuring the Mobility Master IP Address on page 98
- Configuring the Loopback IP Address on page 99
- Configuring Static IP Routes on page 99

Click the following links to configure advanced WLAN or optional WLAN features:

- GRE Tunnels on page 100
- GRE Tunnel Groups on page 107
- Jumbo Frame Support on page 108
- PVST+ on page 110
- RSTP on page 111
- PortFast and BPDU Guard for Spanning Tree on page 113
- LLDP on page 115

Campus WLAN Workflow

Create a campus WLAN by using the new WLAN wizard in the WebUI, manually configuring the WLAN in the WebUI, or manually configuring the WLAN in the CLI.

Using the New WLAN Wizard in the WebUI

The simplest way to create a new WLAN is to use the **New WLAN** wizard, available in the **Configuration > WLANs** section of the WebUI (**Managed Network** node hierarchy). The wizard walks you through the steps to define and configure the SSID, VLAN, authentication and authorization settings, and default user role for the WLAN. The configuration options that appear in the WLAN wizard will vary, depending upon the type of WLAN you choose to create.

Manually Configuring the WLAN in the WebUI

The following workflow lists the tasks to configure a campus WLAN, with a signal SSID, that uses 802.1X authentication. Click any of the links below for details on the configuration procedures for that task.

1. Configure your authentication servers.
2. Create an authentication server group and assign the authentication servers you configured in step 1 to that server group.
3. Configure a firewall access policy.
4. Create a user role and assign the firewall access policy you created in step 3 to that user role.
5. Create an AAA profile.
   a. Assign the user role defined in step 4 to the **802.1X Authentication Default Role** of the AAA profile.
   b. Associate the server group you created in step 2 to the AAA profile.
6. Create a new SSID profile.
7. Create a new virtual AP profile.
8. Associate the virtual AP profile to the AAA profile you created in Step 5.
9. Associate the virtual AP profile to the SSID profile you created in Step 6.

Manually Configuring the WLAN in the CLI

The example below follows the suggested order of steps to configure a virtual AP using the command-line interface.

```
(host) [mynode] (config) #aaa server-group THR-DOT1X-SERVER-GROUP-WPA2
auth-server Internal

(host) [mynode] (config) #ip access-list session THR-POLICY-NAME-WPA2
user any any permit
(host) [mynode] (config) #user-role THR-ROLE-NAME-WPA2
access-list session THR-POLICY-NAME-WPA2

(host) [mynode] (config) #aaa server-group THR-DOT1X-SERVER-GROUP-WPA2
auth-server Internal

(host) [mynode] (config) #aaa profile THR-AAA-PROFILE-WPA2
dot1x-default-role THR-ROLE-NAME-WPA2
dot1x-server-group THR-DOT1X-SERVER-GROUP-WPA2

(host) [mynode] (config) #wlan ssid-profile THR-SSID-PROFILE-WPA2
essid THR-WPA2
opmode wpa2-aes

(host) [mynode] (config) #wlan virtual-ap THR-VIRTUAL-AP-PROFILE-WPA2
ssid-profile THR-SSID-PROFILE-WPA2
aaa-profile THR-AAA-PROFILE-WPA2
vlan 60

(host) [mynode] (config) #ap-group THRHQ1-STANDARD
virtual-ap THR-VIRTUAL-AP-PROFILE-WPA2
```

Understanding VLAN Assignments

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

1. The default VLAN is the VLAN configured for the WLAN.

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

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

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

```
Aruba-User-VLAN
Aruba-Named-User-VLAN
```

The following sections describe:

- VLAN Derivation Priorities for VLAN types on page 82
- Configuring Multiple Wired Uplink Interfaces (Active-Standby) on page 83

VLAN Derivation Priorities for VLAN types

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

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

17. VLAN from VSA during 802.1X
18. VLAN from Microsoft Tunnel attributes during 802.1X
19. VLAN from DHCP options role
20. VLAN from DHCP options



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

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

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

Configuring Multiple Wired Uplink Interfaces (Active-Standby)

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

To allow Mobility Master to obtain a dynamic IP address for a VLAN, enable the DHCP or PPPoE client on Mobility Master for the VLAN. For more information, see Assigning a Static IP Address to a VLAN

Configuring VLANs

Managed Devices operate as layer-2 switches that use a VLAN as a broadcast domain. As a layer-2 switch, the managed device requires an external router to route traffic between VLANs. The managed device can also operate as a layer-3 switch that can route traffic between VLANs defined on Mobility Master.

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

You can optionally configure an IP address and netmask for a VLAN. 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 managed device are forwarded according to the managed device's IP routing table.

The maximum number of VLANs supported on the managed device is 256 each for static VLANs and for dynamic VLANs.



Mobility Masters do not honor CoA to change VLANs. If a VLAN is switched using CoA, clients do not get IP addresses from the new VLAN. Issue the **ip mobile proxy block-dhcp-release** or **aaa user delete** command to get IP address from the new VLAN.

Configuring VLANs include:

- Creating and Updating VLANs on page 84
- Role Derivation for Named VLAN Pools on page 85
- VLAN Pooling Resiliency on page 86
- Adding a Bandwidth Contract to the VLAN on page 87

- Optimizing VLAN Broadcast and Multicast Traffic on page 87
- Inter-VLAN Routing on page 88
- Configuring Source NAT to Dynamic VLAN Address on page 89
- Configuring Source NAT for VLAN Interfaces on page 89

Creating and Updating VLANs

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

Creating and Updating a Single VLAN

The following procedure creates and updates a single VLAN:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces > VLANs** tab.
2. Click **+** to create a new VLAN. (To edit an existing VLAN, click the VLAN entry.) See Creating and Updating VLANs on page 84 to create a range of VLANs.
   a. Enter a name for the new VLAN.
   b. In the **VLAN ID/Range** field, enter a valid VLAN ID. (Valid values are from 1 to 4094, inclusive).
   c. Click **Submit**.
3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.
6. To add physical ports to the VLAN, go to the **Ports** page. To associate the VLAN with specific port-channels, select **Port-Channels**.
   a. If you selected **Port**, select the ports you want to associate with the VLAN from the **Ports** table. For each port, select the new VLAN from the **VLAN** drop-down list.
   b. If you selected **Port-Channel**, select the specific channel number you want to associate with the VLAN from the **Port Channel** table.
7. Click **Submit**.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command creates and updates a single VLAN:

```
(host) [mynode] (config) #vlan <id>
(host) [mynode] (config) #interface gigabitethernet <slot>/<module>/<port>
(host) [mynode] (config-if) #switchport access vlan <vlan>
```

Creating and Updating Bulk VLANs

The following procedure creates and updates bulk VLANs:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces > VLANs** tab.

2. To add multiple VLANs at one time, click **+** on the **VLANs** window.
    a. In the **New VLAN** pop-up window, enter a range of VLANs in the **VLAN ID/Range** field that you want to create at once. For example, to add VLAN IDs numbered 200-300 and 302-350, enter 200-300, 302-350.
    b. Click **Submit**.
3. To add physical ports to a VLAN, select the VLAN.
    a. In the **Port Members** table, click **Edit**.
    b. Select and move the ports from the **Available** list to the **Selected** list.
    c. Click **OK**.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command creates and updates bulk VLANs:

```
(host) [mynode] (config) #vlan <id>
(host) [mynode] (config) #vlan <id> range <range>
```

Creating and Updating a Named VLAN

Refer to the section Address Pool Management on page 259.

Role Derivation for Named VLAN Pools

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



A VLAN name cannot be modified.

For tunnel mode, named VLANs that have the assignment type **hash** and **even** are supported.

For bridge mode only, named VLANs with the assignment type **hash** are supported. If a named VLAN with **even** assignment is assigned to a user rule, user role, server derivation or VSA, then the **hash** assignment is applied and the following error message is displayed - **named VLAN assignment type EVEN not supported for bridge. Applying HASH algorithm to retrieve vlan-id**.



L2 roaming is not supported with an even VLAN assignment.

The following procedures configure Named VLANs under user rule, server derivation, user derivation, and VSA:

To apply a named VLAN to a user rule:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Authentication > User Rules**.
2. Select a user rule from the **User Rules Summary** table.
3. Click **+** to add a new rule.
4. Select **VLAN** from the **Set Type** drop-down list.
5. Select a VLAN from the **VLAN** drop-down list.

6. Configure the remaining profile settings: **Rule Type**, **Condition**, **Value**, and **Description**. Users are assigned the selected VLAN when the rule matches.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

To apply a named VLAN to a user role:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Roles & Policies > Roles**.
2. Select a role from the **Roles** table, and then click **Show Advanced View**.
3. Under **More**, select a VLAN from the **VLAN** drop-down list.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

To apply a named VLAN to a server derivation (server group):

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > Authentication > Auth Servers**.
2. Select a server group from the **Server Groups** table.
3. Under **Server Rules**, click **+** to add a new rule.
4. Select **set vlan** from the **Action** drop-down list.
5. Select a VLAN from the **Vlan** drop-down list.
6. Configure the remaining profile settings: **Attribute**, **Operation**, and **Operand**. Users are assigned the selected VLAN when the rule matches.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command applies a named VLAN in a user rule:

```
(host) [mynode] (config) #aaa derivation-rules user <name>
(host) [mynode] (user-rule) #set vlan condition <rule-type> <attribute> <value> set-value {<role>|<vlan>} [description <rule description>][position <number>]
```

The following CLI command applies a named VLAN in a user role:

```
(host) [mynode] (config) #user-role <name>
(user) [mynode] (config-role) #vlan <string>
```

The following CLI command applies a named VLAN in server derivation:

```
(host) [mynode] (config) #aaa server-group <group>
(user) [mynode] (Server Group) set vlan condition <attribute> contains|ends-with|equals|not-equals|starts-with <string> set-value <set-value-str> [position <number>]
```

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

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

VLAN Pooling Resiliency

Starting from ArubaOS 8.7.0.0, the VLAN pool resiliency feature automatically assigns clients to the next available VLAN ID if a particular VLAN pool is full. The following CLI commands configure VLAN pooling resiliency:

```
(host) [mynode] (config) #vlan-name <name> assignment even ip-timeout
(host) [mynode] (config) #vlan-name <name> assignment even max-ip-timeouts
(host) [mynode] (config) #vlan-name <name> assignment even full-period
```

Users can enable or disable this feature using the **ip-timeout** parameter. The **ip-timeout** parameter configures the timeout value (in seconds) of a DHCP request. If the client is not assigned an IP address within the stipulated time, the request is timed out. After three consecutive IP timeouts, the VLAN ID will be marked as full and clients will be assigned to the next available VLAN ID. The default value for the maximum number of IP timeouts is 3 and it can be configured or edited using the **max-ip-timeouts** parameter. The time duration for which a VLAN ID is marked as full is configured using the **full-period** parameter.

```
This feature is enabled by default and takes effect only when the VLAN assignment type
is EVEN.
```

Adding a Bandwidth Contract to the VLAN

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

The command in the example below adds the MAC address for CDP and VTP to the list of protocols that are not limited by VLAN bandwidth contracts.

```
(host) [mynode] (config) #vlan-bwcontract-explist mac <mac>
```

To show entries in the VLAN bandwidth contracts MAC exception list execute the following command:

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

Optimizing VLAN Broadcast and Multicast Traffic

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

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



If you enable broadcast-multicast optimization on uplink ports, the managed device-generated Layer-2 packets will be dropped.

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

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

The following procedure configures broadcast-multicast optimization:

1. In the **Mobility Master**node hierarchy, navigate to **Configuration > Interfaces > VLANs**.
2. Select a VLAN from the **VLANs** table.
3. Under **Vlan Ids**, select the VLAN ID number.
4. Navigate to the **IPv4** tab for the selected VLAN ID.
5. Click **Other Option** to expand it.
6. Set the **BCMC optimization** to **Enabled** for the selected VLAN.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures broadcast-multicast optimization:

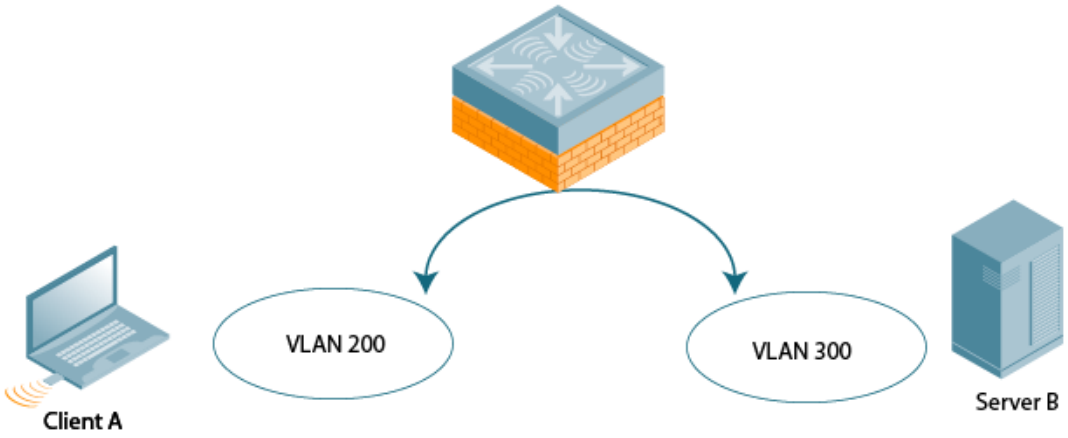
```
(host) [mynode] (config) #interface vlan <vlan>
(host) [mynode] (config-subif)#bcmc-optimization
(host) [mynode] (config-subif)#show interface vlan <vlan>
```

Inter-VLAN Routing

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

In Figure 10, 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 managed device to prevent the flow of traffic between the VLANs.

Figure 10 *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 managed device. Forwarding of inter-VLAN traffic is blocked.
- IP mobility does not work when a mobile client roams to the restricted VLAN. You must ensure that a mobile client on a restricted VLAN is not allowed to roam to a non-restricted VLAN. For example, a mobile client on a guest VLAN will not be able to roam to a corporate VLAN.

The following procedure disables Layer-3 forwarding for a VLAN configured on the managed device:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Interfaces > VLANs** tab.
2. Select a VLAN from the **VLANs** table.
3. Under **Vlan Ids**, select the VLAN ID number.
4. Navigate to the **IPv4** tab for the selected VLAN ID.
5. Click **IP Address Assignment** to expand it.
6. In the **IP assignment** field, configure the VLAN to either obtain an IP address dynamically (via DHCP or PPPoE) or to use a static IP address and netmask.
7. Click **Submit**.
8. Click **Other Option** to expand it.
9. Set **Inter-VLAN routing** to **Disabled**.
10. Click **Submit**.
11. Click **Pending Changes**.
12. In the **Pending Changes** window, select the check box and click **Deploy changes**.

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 addresses 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.

You can configure the source NAT to dynamic VLAN address using the WebUI or CLI.

In the following example, the rule for a guest policy denies traffic to any network addresses.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Roles & Policies > Policies** page.
2. Click **+** to add the policy **guest** with the session type **MAC**.
3. Select the new **guest** policy from the **Policies** table.
4. To add a rule, click **+** in the **Policies > guest** table.
    a. Select **Deny** from the **Action** drop-down list.
    b. Select **Any** from the **MAC address** drop-down list.
    c. Set **Mirror** to **Disabled**.
    d. Click **Submit**.
    e. Click **Pending Changes**.
    f. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command specifies the rule for a guest policy that denies traffic to internal network addresses:

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

Configuring Source NAT for VLAN Interfaces

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

All outbound traffic can enable NAT with the IP address of the VLAN interface as the source address; while the locally routed traffic is sent without any address translation.

IP NAT Inside and IP NAT Outside

In ArubaOS, IP NAT Inside feature allows the user traffic to perform NAT with the desired IP address of the managed device VLAN as the source address. Hence, no new routes need to be added to the existing wired network for the user VLAN.

Traditionally, ArubaOS supported only IP NAT Inside feature that was useful for only traffic going out of uplink VLAN interface. However, the traffic that needed local routing was also going through unnecessary address translation.

IP NAT Outside feature changes the source IP of all the packets coming from downstream network to the IP address of the egress VLAN interface where IP NAT Outside is configured. This feature allows only outbound traffic to perform NAT.

Important Points to Note

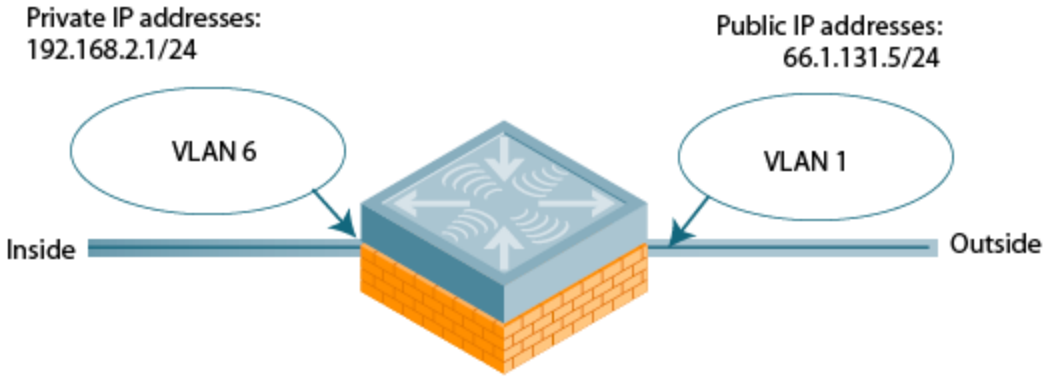
- All ports on the managed device are assigned to VLAN 1 by default. Do not enable the **IP NAT Inside** option for VLAN 1, as this prevents IPsec connectivity between the managed device and its IPsec peers.
- IP NAT Outside must be configured only on the Egress VLAN interface on the managed device, whereas IP NAT Inside must be configured on each and every VLAN interface where the traffic is routed through source NAT.
- IP NAT Outside and IP NAT Inside follow the same rate limit of 40 kbps.
- IP NAT Outside takes precedence over IP NAT Inside, whereas user-defined ACLs take precedence over IP NAT configuration.

You can configure the source NAT for VLAN interfaces using the WebUI or CLI. Following is a sample configuration.

Sample Configuration

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

Figure 11 *Example: Source NAT using the IP Address of Managed Device*



The following procedure configures source NAT for VLAN interfaces:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Interfaces > VLANs** tab. Click **+** to configure VLAN 6 (VLAN 1 is configured through the Initial Setup).
   a. Enter **6** for the VLAN ID.
   b. Click **Submit**.
2. Select VLAN 6.
   a. Navigate to the **IPv4** tab for VLAN 6.
   b. Click **IP Address Assignment** to expand it.
   c. Enter **192.168.2.1** for the **IP address**.
   d. Click **Submit**.
   e. Click **Other Option** to expand it.
   f. Set **NAT Inside** to **Enabled**.
   g. Click **Submit**.
   h. Click **Pending Changes**.
   i. In the **Pending Changes** window, select the check box and click **Deploy changes**.
3. Select VLAN 1.
   a. Navigate to the **IPv4** tab for VLAN 1.
   b. Click **IP Address Assignment** to expand it.
   c. Enter **66.1.131.5** for the **IPv4 address**.
   d. Click **Submit**.
   e. Click **Other Option** to expand it.
   f. Set **NAT Outside** to **Enabled**.
   g. Click **Submit**.
   h. Click **Pending Changes**.
   i. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures source NAT for VLAN interfaces:

```
(host) [mynode] (config) #interface vlan 1
ip address 66.1.131.5 255.255.255.0
ip nat outside
(host) [mynode] (config) #interface vlan 6
ip address 192.168.2.1 255.255.255.0
ip nat inside
```

Trusted and Untrusted Ports and VLANs

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

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



For more information on configuring trusted and untrusted ports or VLANs, see Configuring Trusted or Untrusted Ports and VLANs on page 96

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.

The following sections describe:

- About Trusted and Untrusted Physical Ports on page 92
- About Trusted and Untrusted VLANs on page 92

About Trusted and Untrusted Physical Ports

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

About Trusted and Untrusted VLANs

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

You can set a range of VLANs as trusted or untrusted in trunk mode. The following table lists the port, VLAN and the trusted or 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 18: *Classifying Trusted and Untrusted Traffic*

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

Assign an IP Address to a VLAN

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

- You can manually assign a static IP address to a VLAN. This is the default method and is described in Assigning a Static IP Address to a VLAN on page 93. At least one VLAN on the managed device must be assigned a static IP address.
- Dynamically assigned from a DHCP or PPPoE server. This is described in Configuring a VLAN to Receive a Dynamic Address on page 93.

Assigning a Static IP Address to a VLAN

You can manually assign a static IP address to a VLAN on the managed device using the WebUI or CLI. At least one VLAN on the managed device should have a static IP address.

The following procedure assigns a static IP address to a VLAN:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Interfaces > VLANs** tab. Select a VLAN from the **VLANs** table, and then select a VLAN ID under **Vlan Ids**.
2. Under **IPv4**, select **Static** from the **IP assignment** drop-down list.
3. Enter the **IPv4 address** of the VLAN interface.
4. Select the **Option-82** check box to forward the circuit-specific information to the DHCP server.
5. Enter an **MTU** value for the VLAN, between 1280 and 1500.
6. Enable or disable **Suppress ARP**. If enabled, the managed device prevents flooding of ARP broadcasts on all untrusted interfaces. This is disabled by default.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command assigns a static IP address to a VLAN:

```
(host) [mynode] (config) #interface vlan <vlan>
ip address <ipaddr> <ipmask>
```

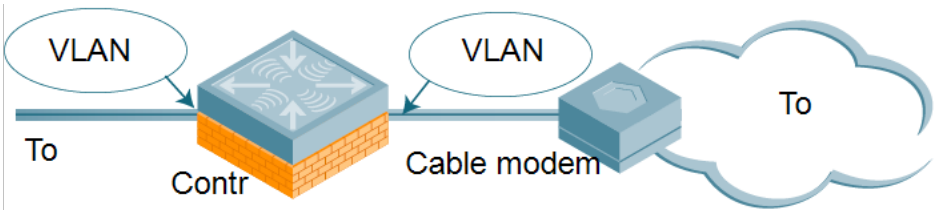
Configuring a VLAN to Receive a Dynamic Address

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

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

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

Figure 12 *IP Address Assignment to VLAN via DHCP or PPPoE*



The following sections describe:

- Enabling the DHCP Client on page 94
- Enabling the PPPoE Client on page 94
- Support for Multiple PPPoE Uplinks on page 95
- Default Gateway from DHCP or PPPoE on page 95
- Configuring DNS or WINS Server from DHCP or PPPoE on page 96

Enabling the DHCP Client

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

The following procedure enables the DHCP client:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Interfaces > VLANs** tab.
2. Select a VLAN from the **VLANs** table, and then select a VLAN ID under **Vlan Ids**.
3. Under **IPv4**, select **DHCP** from the **IP assignment** drop-down list.
4. Enter the **Client ID** and select a link from the **Uplink wired** drop-down list.
5. Enter a priority value for the VLAN ID in the **UpLink Priority** field. All wired uplink interfaces have the same priority by default. If you want to use an active-standby topology, then prioritize each uplink interface by entering a different priority value (1– 4) for each uplink interface.
6. Enter the **Uplink weight** and **MTU** value for the VLAN, between 1280 and 1500.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

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

```
(host) [mynode] (config) #interface vlan 62
(host) [mynode] (config) #uplink wired vlan 62 priority 2
(host) [mynode] (config) #interface vlan 62 ip address dhcp-client client-id myclient
```

Enabling the PPPoE Client

To authenticate the BRAS and request a dynamic IP address, the managed device 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.

The following procedure enables the DHCP client:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Interfaces > VLANs** tab.
2. Select the previously-created VLAN.
3. Select a VLAN from the **VLANs** table, and then select a VLAN ID under **Vlan Ids**.
4. Under **IPv4**, select **PPPoE** from the **IP assignment** drop-down list.

5. Enter the **Service name**, **User name**, and **Password** for the PPPoE session.
6. Enter a priority value for the VLAN ID in the **UpLink Priority** field. All wired uplink interfaces have the same priority by default. If you want to use an active-standby topology, then prioritize each uplink interfaces by entering a different priority value (1– 4) for each uplink interface.
7. Enter an **MTU** value for the VLAN, between 1280 and 1500.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

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

(host) [mynode] (config) #`interface vlan 14`

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

Support for Multiple PPPoE Uplinks

When the same gateway IP address is assigned to multiple PPPoE links, the managed device is unable to install multiple default routes with the same next-hop address. This results in routing issues and leads to health check failure for the PPPoE links. To address this issue, starting ArubaOS 8.4.0.0, a managed device can be configured to support the same gateway IP address over multiple PPPoE uplinks. You can configure the PPPoE default rroute in the datapath using the CLI.

The following CLI command displays the IP address used to configure PPPoE default route in the datapath:

```
(host) [mynode] (config) #interface vlan 5
ip pppoe-username <username>
ip pppoe-password <password>
ip pppoe-gateway-nat 192.168.1.2
!
(host) [mynode] (config) #show ip pppoe-info
pppoe-username <username>
ip pppoe-password <password>
ip pppoe-service-name <service_name>
pppoe vlan: 3
gateway nat: enabled     IP:192.168.1.2
```

Default Gateway from DHCP or PPPoE

The following procedure specifies that the router IP address obtained from the DHCP or PPPoE server be used as the default gateway for the managed devices:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Interfaces > IP Routes** tab.
2. Click **Dynamic Default Gateway** to expand it. Select the following check boxes:
   - **DHCP** - Use DHCP when available to obtain default gateway.
   - **PPPoE** - Use PPPOE when available to obtain default gateway.
   - **Cellular** - Use Cell interface when available to obtain default gateway.
3. Click **Submit**.
4. Click **Pending Changes**.

5. In the **Pending Changes** window, select the check box and click **Deploy changes**.
6. The following CLI command specifies the router IP that can as the default gateway for the managed devices:

```
(host) [mynode] (config) #ip default-gateway import {cell|cell-cost
<cost>|dhcp|dhcp cost <cost>|pppoe|pppoe cost <cost>}
```

Configuring DNS or WINS Server from DHCP or 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 managed device's internal DHCP server.

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

The following procedure configures the DNS or WINS server from DHCP or PPPoE:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > DHCP** tab.
2. Enable the required **DHCP server** - **IPv4** or **IPv6**.
3. Under **Pool Configuration**, click **+**.
4. In the **Add New Pool Configuration** window appears.
5. For **Pool Name**, enter employee-pool.
6. For **Default Routers**, enter 10.1.1.254.
7. For **DNS Servers**, enable **Import from DHCP/PPPoE**.
8. For **WINS Servers,** enable **Import from DHCP/PPPoE**.
9. For **Network**, enter 10.1.1.0 for **IP address** and 255.255.255.0 for **IP mask**.
10. Click **Submit**.
11. Click **Pending Changes**.
12. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures the DNS or WINS server from DHCP or PPPoE:

```
(host) [mynode] (config) #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 Trusted or Untrusted Ports and VLANs

You can configure an Ethernet port as an untrusted access port or configure trusted and untrusted ports and VLANs in trunk mode. Use the following procedures to define access ports and VLANs as trusted or untrusted. For more information on trusted vs untrusted ports and VLANs, see Trusted and Untrusted Ports and VLANs on page 91

Configuring an Ethernet port as an Untrusted Access Port

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

The following procedure configures an Ethernet port as an untrusted access port:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Interfaces > Ports** tab.
2. Select the port you want to configure from the **Ports** table.
3. Select the **Trust** check box to make the port trusted. The default is **Untrusted**.
4. In the **Mode** drop-down list, select **Access**.
5. From the **VLAN** drop-down list, select the **VLAN** whose traffic will be carried by this port.
6. Select the **VLAN trust** check box to make the VLAN trusted. The default is **Untrusted**.
7. In the **VLAN 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. Select whether **Tunneled node** should be **Enabled** or **Disabled**.
9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure an Ethernet port as an untrusted access port:

```
(host) [mynode] (config) #interface range gigabitethernet <slot>/<module-start>/<port-
start>-<module-end>/<port-end>
(host) [mynode] (config-if)#switchport access
(host) [mynode] (config-if)#no trusted
(host) [mynode] (config-if)#switchport access vlan <vlan>
(host) [mynode] (config-if)#no trusted vlan <vlan>
(host) [mynode] (config-if)#ip access-group ap-acl session vlan <vlan>
(host) [mynode] (config-if)#ip access-group validuserethacl in
(host) [mynode] (config-if)#ip access-group validuserethacl out
(host) [mynode] (config-if)#ip access-group validuser session
```

Configuring Trusted and Untrusted Ports and VLANs in Trunk Mode

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

The following procedure configures trusted and untrusted ports and VLANs in trunk mode:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Interfaces > Ports** tab.
2. Select the port you want to configure from the **Ports** table.
3. For **Mode** select **Trunk**.
4. To specify the native VLAN, select a VLAN from the **Native VLAN** drop-down list.
5. Choose one of the following options from the **Allowed VLANs** drop-down list to control the type of traffic the port carries:
   - **Allow all**: The port carries traffic for all VLANs.
   - **Allow specified VLANs**: The port carries traffic for all VLANs selected. Click + to specify a **VLAN**. You can select whether the VLAN is **Trusted** or **Untrusted**.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure trusted and untrusted ports and VLANs in trunk mode:

```
(host) [mynode] (config) #interface gigabitethernet <slot>/<module>/<port>
(host) [mynode] (config-if)#description <string>
(host) [mynode] (config-if)#trusted {vlan <word>}
```

```
(host) [mynode] (config-range)#switchport mode trunk
(host) [mynode] (config-if)#switchport trunk native vlan <vlan>
(host) [mynode] (config-range)#ip access-group test session vlan <vlan>
```

Configuring the Mobility Master IP Address

The IP address of the Mobility Master or managed device is used to communicate with external devices such as APs.



IP addresses used by the Mobility Master or managed device are not limited to its own IP address.

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

If the IP address is not configured, the IP address of the Mobility Master or managed device 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 IP address. You can configure the Mobility Master IP address using the WebUI or CLI.

The following procedure configures the Mobility Master IP address:

1. In the **Mobility Master** or **Managed Network** node hierarchy, select the device and navigate to the **Configuration > System > General** page.
2. Expand the **Controller IP address** section.
3. Select the address you want to set as the IP address of the Mobility Master or managed device from the **IPv4 address or IPv6 address** drop-down lists. This list only contains VLAN IDs with statically assigned IP addresses. If you have previously configured a loopback interface IP address, then it will also appear in this list. Dynamically assigned IP addresses, such as DHCP or PPPOE, do not appear.

   In a native IPv6 deployment scenario, the configuration of IPv4 address is not mandatory. Hence, you can select **None** from the **IPv4 address** drop-down list to remove the IPv4 address of the Mobility Master, when IPv6 address is configured on the Mobility Master.

4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

   Any change in the IP address of the Mobility Master or managed device requires a reboot.

7. In the **Mobility Master** node hierarchy, select the device and navigate to the **Maintenance > Software Management > Reboot** page to reboot Mobility Master and apply the IP address update.
8. Click **Yes** to save the configuration.
9. Click **Reboot**.
10. **Mobility Master** boots up with the updated IP address of the selected VLAN ID.

The following CLI command configures the Mobility Master IPv4 address:

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

The following CLI command removes the IPv4 address of the Mobility Master, when a valid IPv6 address is configured during migration of pure IPv4 or dual-stack deployment to native IPv6 deployment:

```
(host) [mynode] (config) #no controller-ip
```

Configuring the Loopback IP Address

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

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

The following procedure configures the loopback IP address:

1. In the **Mobility Master** or **Managed Network** node hierarchy, select the device and navigate to the **Configuration > System > General** tab.
2. Expand the **Loopback Interface** section.
3. Enter an address into the **IPv4 Address** or **IPv6 Address** field, as required.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.



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

7. In the **Mobility Master**node hierarchy, select the device and navigate to the **Maintenance > Software Management > Reboot** page to reboot Mobility Master and apply the loopback IP address update.
8. Click **Yes** to save the configuration.
9. Click **Reboot**.
10. **Mobility Master** boots up with the changed loopback IP address.

The following CLI command configures the loopback IP address:

```
(host) [mynode] (config) #interface loopback ip address <ipaddr>
```

Configuring Static IP Routes

The following procedure configures a static IP route (like a default route) on a Mobility Master:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Interfaces > IP Routes** tab.
2. Expand the **IP Routes** section.
3. Click **+** to add a static route to a destination network or host.
4. Select the **IP Version**.
5. Enter the **Destination IP address** and **Destination network mask** (255.255.255.255 for a host route)
6. Select a forwarding setting:
   - **Using Forwarding Router Address**: Enter the nexthop IP address in dotted decimal format (A.B.C.D). Optionally, enter the distance metric (cost) for this route. The cost prioritizes routing to the destination. The lower the cost, the higher the priority.
   - **Using Null Interface**: Designate a null interface.
7. Enter the **Next hop IP address** and **Cost**.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures a static IP route on a Mobility Master:

```
(host) [mynode] (config) #ip route <destip>  <destmask> {ipsec <name> [<cost>]|null
<0-0>|<nexthop> [<cost>]}
```

GRE Tunnels

Mobility Master supports GRE tunnels between managed device and other network devices that support GRE tunnels. This section contains the following information:

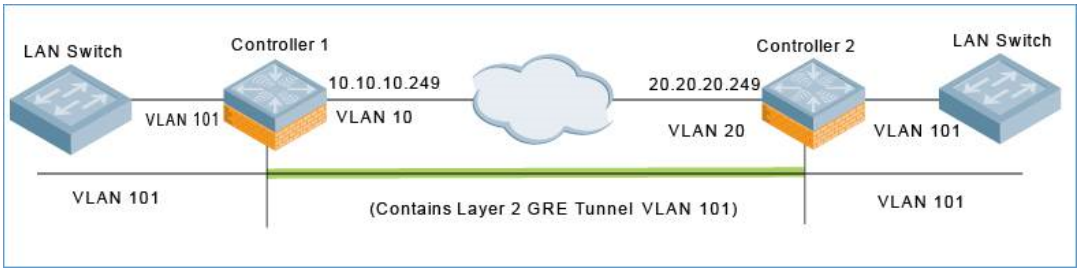
- Layer-2 GRE Tunnels
- Layer-3 GRE Tunnels
- Directing Traffic into the GRE Tunnel
- Configuring Tunnel Keepalives

Layer-2 GRE Tunnels

Layer-2 GRE tunnels allow you to have the same VLAN in multiple locations (separated by a Layer-3 network) and be connected. The forwarding method for a Layer-2 GRE tunnel is bridging.

However, the drawback of using Layer-2 GRE tunnels is that all broadcasts are flooded through the tunnel, adding traffic load to the network and the managed devices. Starting from ArubaOS 8.4.0.0, both trusted and untrusted VLANs are supported on a single Layer-2 GRE tunnel.

Figure 13 *Layer-2 GRE Tunnel*



The traffic flow illustrated by Figure 13 is as follows:

1. The frame enters the source managed device (Controller-1) on VLAN 101.
2. The frame is bridged through Controller-1 into the Layer-2 GRE tunnel.
3. The frame is encapsulated in a GRE packet.
4. The GRE packet enters the network on VLAN 10, is routed across the network to the destination managed device (Controller-2), and then exits the network on VLAN 20.
5. The source IP address of the GRE packet is the IP address of the interface in VLAN 10 in Controller 1.
6. The frame is de-encapsulated and bridged out of the destination managed device (Controller-2) on VLAN 101.

The following procedure configures a Layer-2 GRE tunnel for a source managed device and destination managed device:

1. In the **Managed Network** node hierarchy, navigate to **Configuration** > **Interfaces** > **GRE Tunnels**.
2. Create a new GRE tunnel by clicking **+** below the **GRE Tunnel** table, or edit an existing GRE tunnel by selecting an entry from the **GRE Tunnel** table.
3. Enter the corresponding GRE tunnel values for this managed device.
4. (Optional) Select the **Enable keepalive** checkbox to enable tunnel keepalive heartbeats. For more information on this feature, see Layer-2 GRE Tunnels on page 100
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.
8. Next, access the destination managed device and navigate to **Configuration** > **Interfaces** > **GRE Tunnels**.
9. Select the tunnel ID of interest from the **GRE Tunnel** table.
10. Use the edit screen to configure the destination managed device.
11. (Optional) Select the **Enable keepalive** checkbox to enable tunnel keepalive heartbeats.
12. Click **Submit**.
13. Click **Pending Changes**.
14. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Referring to Figure 13, the following are the required configurations to create the Layer-2 GRE tunnel between controllers named Controller-1 and Controller-2:

IPv4 Controller-1 Configuration

```
(host) [mynode] (config) # interface tunnel 101
description "IPv4 Layer-2 GRE 101"
tunnel mode gre 1
tunnel source vlan 101
tunnel destination 192.168.1.1
tunnel keepalive
trusted
tunnel vlan 101
trusted vlan 101
```

IPv4 Controller-2 Configuration

```
(host) [mynode] (config) # interface tunnel 201
description "IPv4 Layer-2 GRE 201"
```

```
tunnel mode gre 1
tunnel source vlan 201
tunnel destination 192.168.2.1
tunnel keepalive
trusted
tunnel vlan 201
trusted vlan 201
```

The following command example configures a Layer-2 GRE tunnel for IPv6:

IPv6 Controller-1 Configuration

```
(host) [mynode] (config) # interface tunnel 301
description "IPv6 Layer-2 GRE 301"
tunnel mode gre 1
tunnel source ipv6 vlan 301
tunel destination ipv6 2001:1:2:2020::1
tunnel keepalive
trusted
tunnel vlan 301
trusted vlan 301
```

IPv6 Controller-2 Configuration

```
(host) [mynode] (config) # interface tunnel 401
description "IPv6 Layer-2 GRE 401"
tunnel mode gre 2
tunnel source ipv6 vlan 401
tunnel destination ipv6 2001:1:2:1010::1
tunnel keepalive
trusted
tunnel vlan 401
trusted vlan 401
```

Layer-3 GRE Tunnels

The benefit of Layer-3 GRE tunnels is that broadcasts are not flooded through the tunnel, so there is less wasted bandwidth and less load on the managed devices. The forwarding method for a Layer-3 GRE tunnel is routing. By default, GRE tunnels are in IPv4 Layer-3 mode.

Figure 14 *IPv4 Layer-3 GRE Tunnel*

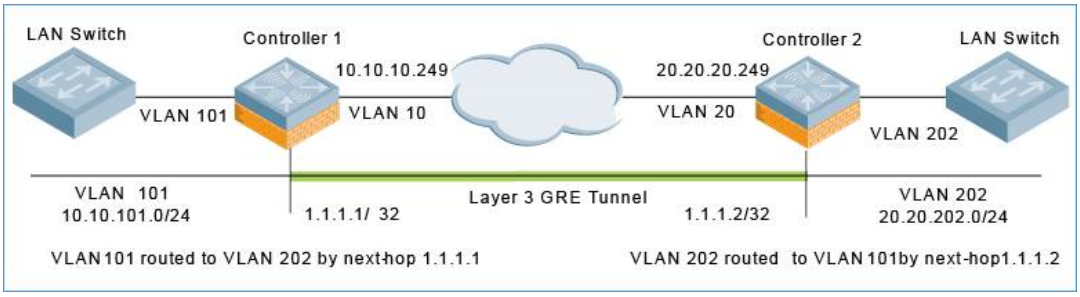
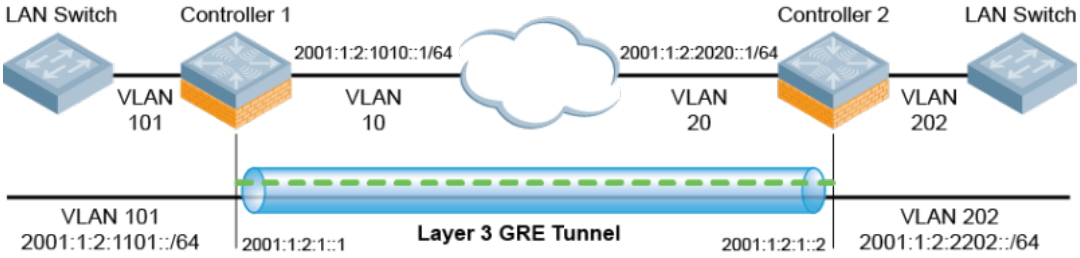


Figure 15 *IPv6 Layer-3 GRE Tunnel*





IPv6 encapsulated in IPv4 and IPv4 encapsulated in IPv6 are not supported. The only Layer-3 GRE modes supported are IPv4 encapsulated in IPv4 and IPv6 encapsulated in IPv6.

The following sections describe:

- Layer-3 Tunnel Traffic Flow on page 103
- Limitations for Static IPv6 Layer-3 Tunnels on page 103
- Layer-3 GRE Tunnels on page 102

Layer-3 Tunnel Traffic Flow

The traffic flow illustrated by and is as follows:

1. The frame enters the source managed device (Controller-1) on VLAN 101.
2. The IP packet within the frame is routed through Controller-1 into the Layer-3 GRE tunnel.
3. The IP packet is encapsulated in a GRE packet.
4. The GRE packet enters the network on VLAN 10, is routed across the network to destination managed device (Controller-2), and then exits the network on VLAN 20.
5. The source IP address of the GRE packet is the IP address of the interface in VLAN 10 in Controller 1.
6. The IP packet is de-encapsulated and routed out of the destination managed device (Controller-2) on VLAN 202.

Limitations for Static IPv6 Layer-3 Tunnels

ArubaOS does not support the following functions for static IPv6 Layer-3 GRE tunnels:

- IPv6 Auto-configuration and IPv6 Neighbor Discovery mechanisms do not apply to IPv6 GRE tunnels.
- The tunnel encapsulation limit and MTU discovery options are not supported on IPv6 GRE tunnels.

The following procedure configures an IPv4 Layer-3 GRE tunnel for Controller-1 and Controller-2:

1. In the source **Managed Network** node hierarchy, select Controller-1
2. Navigate to **Configuration** > **Interfaces** > **GRE Tunnels**. The **GRE Tunnels** page is displayed.
3. Create a new GRE tunnel by clicking + below the GRE Tunnel table, or edit an existing GRE tunnel by selecting that entry in the GRE Tunnel table. The **GRE Tunnel** configuration options appear.
4. Click the IP Version drop-down list and select IPv4 or IPv6.
5. Enter the corresponding GRE tunnel values for the controller.
   a. To configure an IPv4 GRE tunnel , use values for Controller-1 based on the network shown in .
   b. To configure an IPv6 GRE tunnel , use values for Controller-1 based on the network shown in .



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

1. (Optional for IPv4 or IPv6 GRE Tunnels) Select **Enable keepalive** to enable tunnel keepalive heartbeats. For more information on this feature, see Layer-3 GRE Tunnels on page 102
2. Click **Submit**.
3. Click **Pending Changes**.

4. In the **Pending Changes** window, select the check box and click **Deploy changes**
5. Next, log into Controller-2 and navigate to **Configuration > Interfaces > GRE Tunnels.**
6. Create a new GRE tunnel by clicking + below the GRE Tunnel table, or edit an existing GRE tunnel by selecting that entry in the GRE Tunnel table. The **GRE Tunnel** configuration options appear.
7. Enter the corresponding GRE tunnel values for this controller.
   a. To create an IPv4 L3 GRE tunnel, use the values for Controller-2 as shown in .
   b. To create an IPv6 L3 GRE tunnel, use the values for Controller-2 as shown in .
8. (Optional for IPv4 or IPv6 GRE Tunnels) Select **Enable keepalive** to enable tunnel keepalive heartbeats.
9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command examples configure an IPv4 Layer-3 GRE tunnel for IPv4 between two controllers:

Referring to , the following are the required configurations to create the IPv4 Layer-3 GRE tunnel between controllers named Controller-1 and Controller-2:

IPv4 Controller-1 Configuration

```
(host) [mynode] (config) # interface tunnel 104
description "IPv4 L3 GRE 104"
trusted
tunnel
mode gre ip
ip address 1.1.1.1 255.255.255.255
source vlan 10
destination 20.20.20.249
```

IPv4 Controller-2 Configuration

```
(host) [mynode] (config) # interface tunnel 204
description "IPv4 L3 GRE 204"
trusted
tunnel
mode gre ip
ip address 1.1.1.2 255.255.255.255
source vlan 20
destination 10.10.10.249
```

The following command example configures a Layer-3 GRE tunnel for IPv6:

IPv6 Controller-1 Configuration

```
(host) [mynode] (config) # interface tunnel 106
description "IPv6 Layer-3 GRE 106"
trusted
tunnel
tunnel mode gre ipv6
ipv6 address 2001:1:2:1::1
tunnel source ipv6 vlan 10
tunnel destination ipv6 2001:1:2:2020::1
```

IPv6 Controller-2 Configuration

```
(host) [mynode] (config) # interface tunnel 206
description "IPv6 Layer-3 GRE 206"
trusted
tunnel
```

```
tunnel mode gre ipv6
ipv6 address 2001:1:2:1::2
tunnel source ipv6 vlan 20
tunnel destination ipv6 2001:1:2:1010::1
```

Directing Traffic into the GRE Tunnel

You can direct traffic into a GRE tunnel by configuring a *Static route*, which directs traffic to the IP address of the tunnel, or a *Firewall policy (session-based ACL)*, that redirects traffic to the specified tunnel ID.

The following sections describe:

- About Configuring Static Routes on page 105
- Configuring a Firewall Policy Rule on page 105

About Configuring 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 IP Routes on page 99 for detailed information on how to configure a static route.



While redirecting traffic into a Layer-3 GRE tunnel via a static route, be sure to use the tunnel IP address of the controller as the next-hop, instead of providing the tunnel IP address of the destination controller.

Referring to , the following are examples of the required static route configurations to direct traffic into the IPv4 Layer-3 GRE tunnel. for Controller-1 and Controller-2:

- For the controller named Controller-1:

```
(host) [mynode] (config) # ip route 20.20.202.0 255.255.255.0 1.1.1.1
```

- For the controller named Controller-2:

```
(host) [mynode] (config) # ip route 10.10.101.0 255.255.255.0 1.1.1.2
```



IP routing is enabled by default and should not be disabled under VLAN interferences for GRE to work.

Configuring a Firewall Policy Rule

Traffic redirected by a firewall policy rule is *not* forwarded to a tunnel that is “down” (see the next section, Directing Traffic into the GRE Tunnel, for more information on how GRE tunnel status is determined).

The following procedure directs traffic into a GRE tunnel via a firewall policy:

1. On the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Policies** tab.
2. Create a new firewall policy by clicking + below the **Policies** table. The **Add Policy** popup window appears.
3. Enter the **Policy Name**.
4. For **Policy Type**, specify **Session** (the default).

5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**
8. To create a new policy rule for that policy, select the new policy in the Policies table, then scroll to the **Add** table (below the **Policies** table) section and click +.
    a. Select the **Rule Type** and click **OK**.
    b. Specify the **IP Version**.
    c. For **Action**, select **Permit** or **Deny**.
    d. Configure any additional settings.
9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command directs traffic into a GRE tunnel via a firewall policy (session-based ACL):

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

Configuring Tunnel Keepalives

The controller determines the status of a GRE tunnel by sending periodic keepalive frames on the Layer-2 or Layer-3 GRE tunnel. When you enable tunnel keepalives, the tunnel is considered down when the keepalives fail repeatedly.

If you configure a firewall policy rule to redirect traffic to the tunnel, traffic is not forwarded to the tunnel until it is up. When the tunnel comes up or goes down, an SNMP trap and logging message is generated. The remote endpoint of the tunnel does not need to support the keepalive mechanism.

The controller sends keepalive frames at 60-second intervals by default and retries keepalives up to three times before the tunnel is considered down. You can change the default values of the intervals:

- For the **interval**, specify a value between 1 and 86400 seconds.
- For the **retries**, specify a value between 1 and 30.
- To interoperate with Cisco network devices, use the **cisco** option.

The following procedure configures keepalives (Heartbeats):

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces > GRE Tunnels** tab.
2. Locate the tunnel ID for which you are enabling keepalives, and select it. The Edit GRE Tunnel screen appears.
3. To enable tunnel keepalives and display the **Heartbeat interval (seconds)** and **Heartbeat Retries** fields, select **Enable keepalive** toggle switch.
    a. Specify a value for **Heartbeat interval (seconds)**.The allowed value is between 1 and 86400 seconds. The default value is 10 seconds.
    b. Specify a value for **Heartbeat Retries**. The allowed value is between 1 and 30. The default value is 3 retries.
4. Click **Submit**.

5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

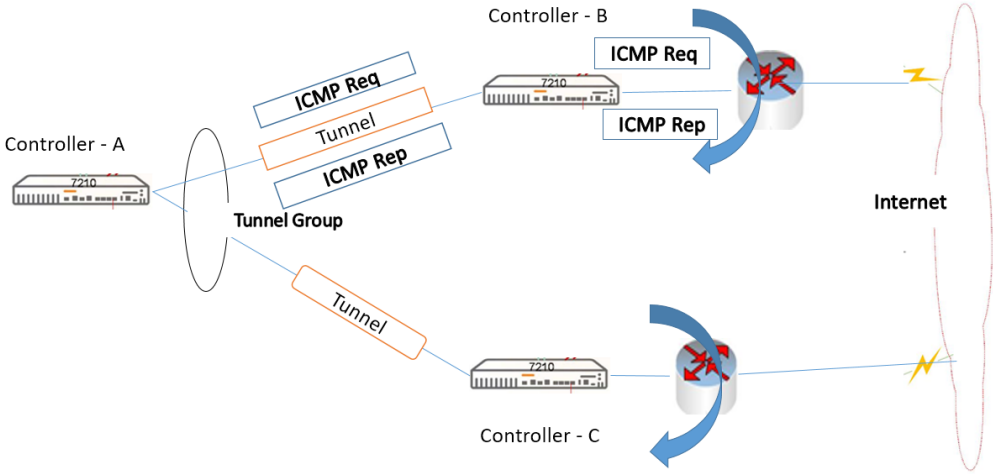
The following CLI command configures the keepalive heartbeats:

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

Configuring ICMP based GRE Tunnels

Starting from ArubaOS 8.5.0.0, GRE tunnel will support ICMP based health-check feature to monitor the status of WAN reachability from remote uplink. ICMP echo requests are periodically sent through the tunnel to a user configured destination. For example in Figure 1, Controller A will send an ICMP echo request to Controller B to ping the destination in WAN . If Controller A does not receive ICMP echo response , it will bring down the tunnel to Controller B and the standby tunnel to Controller C will become active. This feature helps in detecting WAN / Internet failure and will signal the controller to pass the traffic through the secondary / standby GRE tunnel.

Figure 16 *ICMP based GRE Tunnel*



The following CLI command configures ICMP-based tunnel keepalives:

```
(host)[mynode](config-submode) #tunnel keepalive icmp <ipaddr> <next-hop>
```

NOTE

<next-hop> parameter must be configured for L2 tunnels only.

GRE Tunnel Groups

ArubaOS supports redundancy of GRE tunnels for both Layer-2 and Layer-3 GRE tunnels. This feature enables automatic redirection of the user traffic to a standby tunnel when the primary tunnel goes down.

A tunnel group is identified by a name or number. You can add multiple tunnels to a tunnel group. The order of the tunnels defined in the tunnel-group configuration specifies their standby precedence. The first member of the tunnel-group is the *primary tunnel*.

A GRE tunnel group combines two tunnels created on a managed device, where one tunnel is active and the other tunnel is the standby. Traffic forwarding can occur on the active tunnel, and the standby tunnel can become active once the active tunnel is down. When the first tunnel fails, the second tunnel

carries the traffic. The third tunnel in the tunnel-group takes over if the second tunnel also fails. In the meantime, if the first tunnel comes up, it becomes the most eligible standby tunnel.

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

When creating a tunnel group, keep in mind the following:

- When a tunnel is added to the tunnel group, the tunnel is used for data traffic only if it is the active tunnel in the group.
- Standby tunnels do not carry any data traffic. However, all tunnels in the group continue to send and receive keepalive packets.
- Only one type of tunnel can be placed into a tunnel group—either Layer-2 or Layer-3. That is, you cannot have a tunnel group consisting of both Layer-2 and Layer-3 tunnels.
- The default value of tunnel group type is Layer-3.
- All tunnels in a Layer-2 tunnel group must be tunneling the same VLAN.
- A Layer-2 tunnel can only be part of one tunnel group.
- The ArubaOS Layer-2 tunnel-group is not interoperable with other vendors. You must set up Layer-2 tunnel groups between Aruba devices only.
- IPv6 tunnel is not supported in tunnel-group. Hence, you cannot add Layer-2 or Layer-3 GRE tunnels to a tunnel-group in both dual-stack and native IPv6 deployments.

The following procedure configures a Layer-2 or Layer-3 tunnel group:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Interfaces > GRE Tunnels**.
2. Click **+** below the **Tunnel Group** table.
3. Specify a name for the tunnel-group in the **Tunnel Group Name** field.
4. In the **Tunnel Group Members** text box, click **+** to add one or more tunnel IDs.
5. Select the IDs and click **OK**.
6. To enable preemption, select the **Enable preemptive-failover Mode** check box. This option is enabled by default.
7. In the **Mode** section, identify the tunnel group type as a layer-2 or layer-3 group.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures a Layer-2 or Layer-3 tunnel group:

```
(host) [mynode] (config) tunnel-group <tungrpname>
(host) [mynode] (config-tunnel-group)#
mode {L2|L3)
preemptive-failover
tunnel <tunnel-id>
```

Jumbo Frame Support

Jumbo frames are the data frames that are larger than 1500 bytes and includes the Layer 2 header and frame check sequence. Jumbo frames functionality can be configured on the following and can support up to 9216 bytes of payload:

- Mobility Controllers (7000 Series controllers, 7200 Series controllers, and Mobility Controller Virtual Appliances.
- Mobility Master (Mobility Master Virtual Appliance and hardware Mobility Master)

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



Jumbo frames are not supported on Mobility Controller Virtual Appliances and Mobility Master Hardware Appliances running on Windows Hyper-V.

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

- **Tunnel node**: In a tunneled node deployment, the wired clients connected on the tunneled nodes can send and receive the jumbo frames.
- **L2 or L3 GRE tunnels**: When you establish a GRE tunnel between two managed devices, the clients on one managed device can send and receive jumbo frames from the clients on the other managed device on enabling jumbo frames.
- **Between wired clients**: In a network where clients connect to the managed device with jumbo frames enabled ports can send and receive the jumbo frames.
- **Wi-Fi tunnel**: A Wi-Fi tunnel can support an AMSDU jumbo frame for an AP (the maximum MTU supported is up to 9216 bytes).



AMSDU is not supported on x86 platforms.

The following sections describe:

- Limitations for Jumbo Frame Support on page 109
- Configuring Jumbo Frame Support on page 109
- Viewing Status of Jumbo Frame Support on page 110

Limitations for Jumbo Frame Support

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

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

Configuring Jumbo Frame Support

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

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

```
(host) [mynode] (config) #firewall jumbo mtu <1789-9216>
```



You can configure the MTU value between 1789-9216. The default MTU value is 9216.

- To enable jumbo frame support on a port channel:

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

- To enable jumbo frame support on a port:

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

Viewing Status of Jumbo Frame Support

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

```
(host) [mynode] #show firewall
```

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

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

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

```
(host) [mynode] #show interface port-channel <id>
```

PVST+

PVST+ provides load-balancing of VLANs across multiple ports, resulting in optimal usage of network resources. PVST+ also ensures interoperability with industry-accepted PVST+ protocols.



PVST+ is disabled by default.

The following sections describe:

- PVST+ Interoperability and Best Practices on page 110
- PVST+ on page 110
- PVST+ on page 110

PVST+ Interoperability and Best Practices

The interoperability between RSTP and PVST+ includes:

- When the access port on the managed device and the trunk port terminate on one Layer 2 switch running PVST+, PVST+ will send untagged STP BPDUs on the access port; it also transmits untagged STP BPDUs (in addition to the other PVST+ BPDUs) on the native VLAN trunk port. If the Aruba managed device is the root, it will detect a loop on the native VLAN.



If PVST+ is not on the managed device, best practices recommend disabling RSTP on the Aruba managed device to avoid a looping issue.

- For VLAN load balancing when managed devices are connected to armed mode, the VLAN priorities on two ports and bridge priorities must be configured so that one set of VLANs are active on one link, and the other set of VLANs are active on the other link.
- Supported instances include: 64 on the 7000 Series and 7200 Series controllers.

The following procedure enables PVST+:

1. In the **Mobility Master**node hierarchy, navigate to **Configuration > Interfaces > VLANs** and select a VLAN with one or more active interfaces in the **VLANs** table.
2. In the **Port Members** table, select the **More** option.
3. Expand the **Spanning Tree** section, and enable **PVST+** mode.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command enables PVST+:

PVST+ is disabled by default. Enable PVST+, ensure a VLAN instance is configured, and then configure PVST+.

1. Enable PVST+:

```
(host) [mynode] (config) #spanning-tree mode rapid-pvst
```

2. Configure PVST+ forward time; the following command sets the time VLAN 2 spends in the listening and learning state (3 seconds):

```
(host) [mynode] (config) #spanning-tree vlan 2 forward-time 3
```

3. Configure PVST+ hello time; the following command sets the time VLAN 2 waits to transmit BPDUs to four seconds:

```
(host) [mynode] (config) #spanning-tree vlan 2 hello-time 4
```

4. Configure PVST+ max age; the following command sets the time VLAN 2 waits to receive a hello packet to 30 seconds:

```
(host) [mynode] (config) #spanning-tree vlan 2 max-age 30
```

5. Configure PVST+ priority: the following command sets the VLAN 2 priority to 10, making it more likely to become the root bridge:

```
(host) [mynode] (config) #spanning-tree vlan 2 priority 10
```

6. Configure PVST+ on a range of VLANs using the VLAN IDs (coma separated or hyphen separated):

```
(host) [mynode] (config) #spanning-tree vlan range 2-6,11
```

RSTP

The ArubaOS implementation of RSTP is as specified in 802.1w, with backward compatibility to legacy 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 Aruba managed devices.

The ArubaOS RSTP implementation interoperates with PVST (Per VLAN Spanning Tree 802.1D) and Rapid-PVST (802.1W) implementation on industry-standard routers or switches. Aruba only supports global instances of STP and RSTP. Therefore, the ports on industry-standard routers or switches must be on the default or untagged VLAN for interoperability with Aruba managed devices.

ArubaOS supports RSTP on the following interfaces:

- FastEthernet IEEE 802.3: fastethernet
- Gigabitethernet IEEE 802.3: gigabitethernet
- Port Channel ID: port-channel

Since RSTP is backwards compatible with STP, it is possible to configure both bridges in the same network. However, such mixed networks may not always provide rapid convergence. RSTP provides rapid convergence when interfaces are configured as either:

- **Edge ports:** These are the interfaces or 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, indicating that it should behave otherwise. It does not go through the Listening and Learning states.
- **Point-to-Point links:** These are the interfaces or ports connected directly to neighboring bridges over a point-to-point link. RSTP negotiates with the neighbor bridge for rapid convergence or transition only when the link is point-to-point.

Table 19: *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, RSTP introduces port roles for all the interfaces (see Table 20).

Table 20: *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 RSTP port interface is designated as point-to-point, by default, in the existing port configuration screen. The following procedure configures RSTP:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > Interfaces > Ports**.
2. In the **Ports** table, click the port number for which you want to enable RSTP or STP features.
3. Select the **Show Advanced Options** link at the bottom of the **Ports** tab.

4. Enable **Spanning Tree**.
5. (Optional) Define values for the following Spanning Tree configuration parameters:
   - **Cost**: Defines the RSTP interface path cost. The cost prioritizes routing to the destination. The lower the cost, the higher the priority.
   - **Priority:** Sets the interface's RSTP priority. The supported range is 0–255, and the default value is 128.
   - **Port Fast**: Changes from blocking to forwarding mode, enabling forwarding of traffic from the interface.
   - **Point-to-Point**: Sets the interface as a point-to-point link.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands change the default configurations:

```
(host) [mynode] (config) #interface gigabitethernet <slot>/<module>/<port>
(host) [mynode] (config-if) #spanning-tree
cost <value>
point-to-point
port-priority <value>
portfast
```

The following CLI commands can be used to view settings and troubleshoot RSTP issues:

- The **show spantree** command displays the root and bridge information, verifying that they are correct. The port or interface information (example: state, role, and so on) is also displayed to make sure that the state and role information correspond with each other. For more details and examples on the **show spantree** command, refer to **show spantree** in the *ArubaOS CLI Reference Guide*.
- The **show spanning-tree interface** command (config-if mode) displays Tx or Rx BPDU counters. For example, if a port's role is designated, it only transmit BPDUs but does not receive any. In this case, the Tx counter continues to increase in increments while the Rx counter remains the same. This is reversed when the role of a port is root/alternate/backup. For more details and examples on the **show spanning-tree interface** command, refer to **show spaning-tree** in the *ArubaOS CLI Reference Guide*.

PortFast and BPDU Guard for Spanning Tree

The PortFast and BPDU Guard features enhance network reliability, manageability, and security for Layer-2 STP .

Some devices and local stacks running on systems or workstations are capable of generating potential STP BPDUs that cause DOS attacks. PortFast and BPDU Guard features provide stability and security for network topologies to prevent such attacks, and can be applied either independently or together.

The following sections describe:

- PortFast on page 114
- BPDU Guard on page 114
- Scenarios Supported on PortFast and BPDU Guard on page 114
- Enabling PortFast on a Port on page 115
- Enabling BPDU Guard on a Port on page 115

PortFast

The PortFast feature is introduced to avoid network connectivity issues. These issues are caused by delays in STP enabled ports moving from blocking-state to forwarding-state after transitioning from the listening and learning states. STP enabled ports that are connected to devices such as a single switch, workstation, or a server can access the network only after passing all these STP states. Some applications need to connect to the network immediately, else they will timeout.

Enabling the PortFast feature causes a switch or a trunk port to enter the STP forwarding-state immediately or upon a linkup event, thus bypassing the listening and learning states. The PortFast feature is enabled at a port level, and this port can either be a physical or a logical port. When PortFast feature is enabled on a switch or a trunk port, the port immediately transitions to the STP forwarding state.

Though PortFast is enabled the port still participates in STP. If the port happens to be part of topology that could form a loop, the port eventually transitions into STP blocking mode. PortFast is usually configured on an edge port, which means the port should not receive any STP BPDUs. If the port receives any STP BPDU, it moves back to normal or regular mode and will participate in the listening and learning states.

In most deployments, edge ports are access ports. However, in this scenario there are no restrictions in enabling the PortFast feature. The mode of the port changes from PortFast to non-PortFast when the port receives a STP BPDU. To re-enable this feature on a port, run the **shut** command followed by a **no-shut** command at the interface or port level.



Configuring PortFast on a non-edge port can cause instability to the STP topology.

BPDU Guard

BPDU Guard feature protects the port from receiving STP BPDUs, however the port can transmit STP BPDUs. When a STP BPDU is received on a BPDU Guard enabled port, the port is shutdown and the state of the port changes to **ErrDis** (Error-Disable) state. The port remains in the **ErrDis** state until the port status is manually changed by using the configuration command **shut** followed by a **no-shut** applied on the interface. In most deployments, BPDU Guard feature is configured over the PortFast enabled STP ports, but in this implementation the BPDU Guard feature can be enabled on any of the STP ports, with or without PortFast feature being enabled on these ports.



It is recommended not to enable the BPDU Guard feature on a trunk port that forms the STP topology.

Scenarios Supported on PortFast and BPDU Guard

PortFast and BPDU Guard features are applied at the port or interface level. These features can also be applied in the following scenarios:

- RSTP and PVST modes
- Access and Trunk ports
- Physical and Logical ports

In the global RSTP mode, there is only one RSTP instance running in the entire Mobility Master. If the port that is enabled with PortFast and BPDU Guard receives any STP BPDU, it affects all ports, as the global RSTP runs on a port basis.

In the PVST mode, there can be multiple instances of RSTP running, as they are based per VLAN. Though it is based per VLAN, it will still behave in the same way as it does in the global RSTP mode. For example,

if there are five VLANs and each VLAN has a separate RSTP instance running, then any STP BPDU received on any of these five ports effects all ports.

If an STP BPDU is received from any one of the five RSTP instances running, the port that is enabled with BPDU Guard shuts down and goes to **ErrDis** state. In other words, both PortFast and BPDU Guard features are applied on a port basis for both global RSTP and PVST modes, even though the PVST runs on a per VLAN basis.

Enabling PortFast on a Port

The following procedure enables PortFast on a port:

1. In the **Mobility Master**node hierarchy, select the device and navigate to **Configuration > Interfaces >Ports**.
2. In the **Ports** table, click the port number for which you want to enable PortFast and BPDU Guard.
3. Select the **Show Advanced Options** link at the bottom of the **Ports** tab .
4. Select the **PortFast** check box.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.



It is recommended to enable PortFast only on access port types. However, PortFast can be enabled on the trunk ports by selecting the **Trunk** check box in the WebUI.

Execute the following commands to enable PortFast:

```
(host) [mynode] (config) #interface gigabitinternet <slot>/<module>/<port>
(host) [mynode] (config-if)#spanning-tree portfast
```

Execute the following commands to disable PortFast:

```
(host) [mynode] (config) #interface gigabitinternet <slot>/<module>/<port>
(host) [mynode] (config-if) #no spanning-tree portfast
```

Execute the following command to enable PortFast on trunk ports:

```
(host) [mynode] (config) #interface gigabitethernet <slot>/<module>/<port>
(host) [mynode] (config-if)#spanning-tree portfast trunk
```

Execute the following show command to display the status of the STP ports:

```
(host) [mynode] (config-if) #show spanning-tree interface gigabitethernet
<slot>/<module>/<port>
```

Enabling BPDU Guard on a Port

The following CLI command enables PortFast and BPDU Guard:

```
(host) [mynode] (config) #interface gigabitinternet <slot>/<module>/<port>

(host) [mynode] (config-if)#spanning-tree bpduguard
```

LLDP

LLDP, defined in the IEEE 802.1AB standard, is a Layer 2 protocol that allows network devices to advertise their identity and capabilities on a LAN. ArubaOS supports a simple one-way neighbor

discovery protocol with periodic transmissions of LLDP PDUs, allowing managed devices to advertise identity information and capabilities to other nodes on the network and store the information discovered about the neighbors.

LLDP supported devices use attributes known as TLVs to receive and send information such as configuration information, device capabilities, and device identity to their neighbors. These TLVs contain type, length, and value descriptions, use the destination MAC address 01:80:c2:00:00:0e, and are constrained to a local link. SNMP support is available for LLDP MIBs.

The following sections describe:

- Supported TLVs on page 116
- LLDP-MED on page 116
- Restrictions and Limitations on page 116
- Configuring LLDP on page 117
- Configuring LLDP-MED on page 117
- Monitoring LLDP Data and Configuration on page 117

Supported TLVs

ArubaOS supports the following basic management TLVs, all of which are enabled by default:

- MAC Phy configuration TLV
- Management address TLV
- Maximum frame size TLV
- Port-description TLV
- Port VLAN ID TLV
- System capabilities TLV
- System description TLV
- System name TLV
- VLAN name TLV

LLDP-MED

LLDP-MED is an extension to LLDP that supports interoperability between VoIP devices and other networking clients. LLDP-MED network policy discovery lets end-points and network devices advertise their VLAN IDs (for example, voice VLAN), priority levels, and DSCP values. ArubaOS supports a maximum of eight LLDP -MED Network Policy profiles.

Creating an LLDP MED network policy profile does not apply the configuration to any AP or AP interface or interface group. To apply the LLDP-MED network policy profile, you must associate it to an LLDP profile, then apply that LLDP profile to an AP wired port profile.

You can use the command, **ap lldp med-network-policy-profile** to define an LLDP MED network policy profile that defines DSCP values and L2 priority levels for a voice or video application.



When you use the default LLDP configuration, the **LLDP RX** and **LLDP TX** parameters are disabled. You must explicitly enable them for LLDP to work.

Restrictions and Limitations

- Inventory-management and Location TLVs are not currently supported.
- Aggregation-management and Power-management TLVs are not supported.
- CDP proprietary is not supported.
- The maximum number of neighbors that can be learned on the managed device (including all the per port neighbors) is 250.

Configuring LLDP

Configure LLDP using the following commands on a specific managed device interface. For detailed information on the LLDP commands, refer to the **interface fastethernet | gigabitethernet** , **show lldp** and **show ap lldp** command in the ArubaOS *CLI Reference Guide.*

```
(host) [mynode] (config) #interface gigabitethernet <slot>/<module>/<port>
(host) [mynode] (config-if) #lldp
fast-transmit-counter <1-8>
fast-transmit-interval <1-3600>
med
receive
transmit
transmit-hold <1-100>
transmit-interval <1-3600>
```



If you use the default LLDP configuration, the **transmit** and **Receive** parameters are disabled. You must explicitly enable them for LLDP to work.

Configuring LLDP-MED

When you create an LLDP MED network policy profile, you must associate it to an LLDP profile, then apply that LLDP profile to an AP wired port profile.

The following commands create a LLDP MED network policy profile for streaming video applications and marks streaming video as high-priority traffic.

```
(host) [mynode] (config) #ap lldp med-network-policy-profile vid-stream
(host) [mynode] (AP LLDP-MED Network Policy Profile "vid-stream") #dscp 48

(host) [mynode] (AP LLDP-MED Network Policy Profile "vid-stream") #l2-priority 6
(host) [mynode] (AP LLDP-MED Network Policy Profile "vid-stream") #tagged
(host) [mynode] (AP LLDP-MED Network Policy Profile "vid-stream") #vlan 10
```

Next, the LLDP MED network policy profile is assigned to an LLDP profile, and the LLDP profile is associated with an AP wired-port profile.

```
(host) [mynode] (config) #ap lldp profile video1
(host) [mynode] (AP LLDP Profile "video1") #lldp-med-network-policy-profile vid-stream
(host) [mynode] (config) #ap wired-port-profile corp2

(host) [mynode] (AP wired port profile "corp2") #lldp-profile video1
```

Monitoring LLDP Data and Configuration

The following show commands display aggregate and per-interface information about LLDP configurations and neighborhood data.

Table 21: *LLDP Show Commands*

Command	Description
`show lldp interface` **gigabitethernet <slot>/<module>/<port>**	Displays LLDP information for all interfaces, or include the optional **gigabitethernet <slot>/<module>/<port>** parameters to display LLDP information for a specific interface.
`show lldp neighbor` **gigabitethernet <slot>/<module>/<port>** `[details]`	This command displays information about LLDP peers, including the name of the neighbor, MAC address and the capabilities of the peer to operate as a router, bridge, access point, phone or other network device. Include the optional **gigabitethernet <slot>/<module>/<port>** parameters to display detailed information about LLDP neighbors for a specific interface, or include the **details** parameter to display additional details about the device type, and LLDP-MED, if applicable.
`show lldp statistics` **gigabitethernet <slot>/<module>/<port>**	Displays information about LLDP TLVs sent and received on all interfaces, or include the optional **gigabitethernet <slot>/<module>/<port>** parameters to display LLDP TLV statistics for a specific interface.
`show ap lldp med-network-policy-profile`	Displays a list of LLDP-MED Network Policy profiles, or display the current configuration settings of an individual profile.
`show ap lldp counters`	Shows LLDP counters for a specific AP, or all APs sending or receiving LLDP PDUs.
`show ap lldp [<profile>]`	Displays a list of LLDP-MED Network Policy profiles, or display the current configuration settings of an individual profile.

The ArubaOS implementation of LACP is based on the standards specified in 802.3AD. LACP provides standardized means for exchanging information with partner systems to form a port-channel group (LAG). LACP avoids port channel misconfiguration.

Two devices (actor and partner) exchange LACP data units when forming a LAG. After multiple ports in the system have the same actor system ID, actor key, partner system ID, and partner key, they belong to the same LAG.

The maximum number of supported port-channels is eight. With the introduction of LACP, this number remains the same.

Two LACP configured devices exchange LACP data units to form a LAG. A device is configurable as an active or passive participant. In active mode, the device initiates data units irrespective of the partner state; passive mode devices respond only to the incoming data units 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 CLI Reference Guide*.

LACP data units exchange their corresponding system identifier or priority along with their port key or priority. This information determines the LAG of a port. The LAG for a port is selected based on its keys. The port is placed in that LAG only when its system ID or key and system ID or key of its partner matches the other ports in the LAG (if the group has ports).

Configuring Port Channel LACP

The following procedure describes the steps to add a new port channel using the WebUI:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces > Ports** tab. In the **Port Channel** table, click **+** to open the **New Port Channel** dialog.
2. In the **New Port Channel** dialog, select a port channel ID from the drop-down list and click **Submit**.
3. In the **Port ID** section, select **LACP** from the **Protocol** drop-down list.
4. Select **Active** in **LACP mode**.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box, and click **Deploy changes**.



For information on configuring LACP on 220 Series and 270 Series access points, see Link Aggregation Support on page 672

The following CLI commands configure LACP settings:

1. Enable LACP and configure the per-port specific LACP.

```
(host)[mynode] (config) #interface  gigabitethernet <slot>/<module>/<port>
```

2. Configure LACP group and mode.

```
(host)[mynode] (config-submode)#lacp group <id> mode {active | passive}
```

- group <id> range is 0–7.
- Active mode—the interface is in an active negotiating state. LACP runs on any link that is configured to be in the active state. The port in an active mode also automatically initiates negotiations with other ports by initiating LACP packets.
- Passive mode—the interface is *not* in an active negotiating state. LACP runs on any link that is configured in a passive mode. The port in a passive mode responds to negotiations requests from other ports that are in an active mode. Ports in passive mode respond to LACP packets.



A port in passive mode cannot set up a port channel (LAG group) with another port in passive mode.

3. Set the timeout for the LACP session. The timeout value is the amount of time that a port-channel interface waits for a LACP data unit from the remote system before terminating the LACP session. The default long time-out value is 90 seconds; short is 3 seconds.
   ```
   (host)[mynode] (config-submode)#lacp timeout {long | short}
   ```
4. Set the port priority. The higher the priority value the lower the priority. The range is 1-65535 and the default is 255.
   ```
   (host)[mynode] (config-submode)#lacp port-priority <value>
   ```
5. Save the configuration.
   ```
   (host)[mynode] (config-submode)#write memory
   ```
6. View your LACP neighbor.

   The port uses the group number +1 as the “actor admin key”. All the ports use the long timeout value (90 seconds) by default.
   ```
   (host)[mynode] (config-submode)#show lacp <id> neighbor
   ```
   The port status is displayed as “DOWN” (see the following example) under the following conditions:
   - When a port in a LAG is misconfigured (the partner device is different than the other ports)
   - If the neighbor ages out
   - If a neighbor cannot exchange LACP data units with the partner

   ```
   (host)[mynode] (config-submode)#show lacp <id> internal
   ```

Additionally, you can configure logging levels for LAGM process and debug LACP related issues on the GSM channel by using the following commands:

Use the following command to configure the logging level for LAGM process:

```
(host)[mynode] (config)# logging system process lagm level <category> [subcat
<subcategory>]
```

Use the following command to debug issues related to LACP in the GSM channel:

```
(host) [mode]# show gsm debug channel port_info
```

LACP Best Practices and Exceptions

Following are the best practices and exceptions to keep in mind when configuring LACP on a managed device:

- LACP is disabled by default.
- LACP depends on periodical Tx/Rx of LACP data units. Any failure detected at a port can be removed from the LAG. Failure detection period depends on the configured timeout.

- The maximum LAG supported per system is eight groups; each group can be created statically or through LACP.
- Each LAG can have up to eight member ports.
- The LAG group ID range is 0–7 for both static (port-channel) and LACP groups.
- When a port is added to a LACP LAG, it inherits the properties of the port-channel. Some of these properties are VLAN membership, trunk status, and so on.
- When a port is added to a LACP LAG, the property (example: speed) of the port is compared to the property of the existing port. If there is a mismatch, the command is rejected.
- The LACP commands cannot be configured on a port that is already a member of a static port-channel. Similarly, if the group assigned in the command **lacp group <number>** already contains static port members, the command is rejected.
- The port uses the group number as its actor admin key.
- All ports use the **long** (90 seconds) timeout value by default.
- The output of the command **show interface port-channel <port-channel ID>** indicates if the LAG is created by LACP (dynamic) or static configuration. If the LAG is created through LACP, you cannot add or delete any ports under that port channel. All other commands are allowed.

LACP Sample Configuration

The following sample configuration is for gigabitethernet port/slot 0/1:

```
(host)[mynode] (config) #interface  gigabitethernet 0/0/4
(host)[mynode] (config-submode)#lacp group 1 mode active
(host)[mynode] (config-submode)#lacp timeout long
(host)[mynode] (config-submode)#lacp port-priority 2
(host)[mynode] (config-submode)#write memory

Saving Configuration...
Partial configuration for /mm/mynode
-------------------------------------
Contents of : /flash/config/partial/0/p=sc=mynode.cfg
interface gigabitethernet 0/0/4
lacp group 1 mode active
lacp port-priority 2
lacp timeout long
!
Configuration Saved
```

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

Typically, the IP address assigned on an IPv6 host consists of a 64-bit subnet identifier and a 64-bit interface identifier. IPv6 addresses are represented as eight colon-separated fields of up to four hexadecimal digits each. The following are examples of IPv6 addresses:
`2001:0000:0eab:DEAD:0000:00A0:ABCD:004E`

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

For example, the address, `2001:0000:0eab:dead:0000:00a0:abcd:004e` can also be represented as:

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

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

```
2001:eab::1/128 - single Host
2001:eab::/64 - network
```

Native IPv6 Support

ArubaOS now provides native IPv6 support that allows enterprises to deploy pure IPv6 wireless network with Mobility Masters and managed devices. In a dual-stack deployment, even though both IPv4 and IPv6 addresses are configured on the networking devices, pure IPv6 deployment is not possible since IPv4 addresses are mandatory to communicate with the APs. In a native IPv6 deployment, all the applications and processes running on the managed devices support IPv6 addresses for seamless communication between Mobility Masters and managed devices. Native IPv6 deployment is applicable in the following scenarios between Mobility Masters and managed devices:

- If both Mobility Master and managed device are configured in pure IPv6 deployment (without configuring IPv4 address), the AP comes up only with a valid IPv6 address on the managed device.
- If both Mobility Master and managed device are configured in pure IPv6 deployment (with managed devices configured in dual-stack deployment but Mobility Masters configured without IPv4 address), the AP comes up on the managed device either with IPv4 or IPv6 address regardless of the cluster.

Native IPv6 deployment impacts the following components:

- VLAN Interface—The IPv4 address on a VLAN interface is now optional for Mobility Masters and managed devices. Hence, depending on the type of deployment, you can configure either IPv4 address, IPv6 address, or a combination of both.
- Master IP Address—It is not mandatory to configure IPv4 address of the Mobility Master or managed devices, when IPv6 address is configured. You can now delete IPv4 address of the Mobility Master, when IPv6 address is configured on the Mobility Master. Also, when managed devices connect to Mobility Masters using VPNC, the IPv4 address on VPNC is optional during master IP configuration.

- VRRP IP Address—In a Layer 2 redundancy scenario, it is not mandatory to configure IPv4 address of VRRP.
- LMS IP Address—The LMS IP address in AMON feeds is now populated with both IPv4 and IPv6 addresses, depending on the type and availability of IPv4 or IPv6 address on the Mobility Master and managed devices. Native IPv6 deployment also determines the type of PAPI transport and tunnel used to transmit the AMON messages.
- ZTP through DHCP—In a native IPv6 deployment, the managed devices obtain IPv6 address also through DHCP option-17 and option-16 fields for finding the Mobility Master during ZTP. Hence, IPv4 address of Mobility Masters is not required in DHCP option-17 field.
- Setup dialog—The full setup dialog now provides flexibility to configure only IPv4, IPv6, or a combination of both. If IPv6 address is used to terminate IPsec tunnel, then it is no longer mandatory to configure IPv4 address in Mobility Master IP configuration in the setup dialog.

The mini setup dialog is not recommended in a native IPv6 deployment.



In a native IPv6 deployment, it takes around 30 minutes to initially bring up the APs running ArubaOS 8.7.1.2 or earlier images. For APs running ArubaOS 8.7.1.3 or later images, it takes around 17 minutes to initially bring up the APs. For information on setting the IP preference for an AP, see the **Provisioning AP to Native IPv6** section of the *ArubaOS 8.x IPv6 Deployment Guide*.

Starting from ArubaOS 8.7.0.0, you can delete the controller IPv4 address at the device and group level while migrating from pure IPv4 or dual-stack deployment to native IPv6 deployment by using the **no controller-ip** command. The following changes are introduced as part of the **no controller-ip** command:

1. You can delete the controller IPv4 address in the following scenarios:
   - When a valid controller IPv6 address is available at the same device or group level.
   - When a single IPv4 address is available on the controller.
2. You cannot delete the controller IPv4 address when multiple IPv4 addresses are available. Hence, you must ensure that only the controller IPv4 and its interface address are the last remaining IPv4 entities to be deleted during the migration process. Depending on the scenario, one of the following errors is displayed in the CLI when you issue the **no controller-ip** command:

   ```
   Controller IPv4 cannot be removed. Please configure controller-ipv6 on some other valid Vlan or loopback
   Controller IPv4 cannot be removed. Multiple v4 addresses exist on the controller
   ```

3. An attempt to delete the controller IPv4 address automatically deletes the last remaining IPv4 addresses on the corresponding VLAN or loopback interface by issuing the following commands:
   - For VLAN interface:

     ```
     interface vlan <id> no ip address
     ```

   - For loopback interface:

     ```
     interface loopback no ip address
     ```

4. An attempt to delete the last remaining IPv4 addresses is prevented by the validation code and displays the following error message in the CLI:

   ```
   Controller IPv4 configured with this address. Execute <no controller-ip> command to auto-delete the interface address.
   ```

Supported Applications

ArubaOS supports native IPv6 deployment for the following applications or scenarios:

- The applications in Mobility Masters and managed devices that are connected directly or through VPNC in an IPv6 network.
- The applications in primary and secondary Mobility Masters that are connected in IPv6 Network.
- Remote AP inner IP pool in cluster deployment.
- ClearPass Policy Manager downloadable user role with RADIUS server configured with IPv6 address.
- Communication with server over the following standard protocols:
  - NTP
  - SNMP
  - SCP
  - FTP
  - TFTP
  - RADIUS
- Configuration of ClientMatch in ARM profiles.
- Configuration of **upgrademgr** process between Mobility Masters and managed devices.
- Scheduled deployments in AirMatch.
- WebCC feature to download database for web classification from cloud service.

Important Point to Remember

- Native IPv6 deployment is currently not supported for the following applications or scenarios:
  - AirGroup
  - UCC
  - Activate for ZTP or whitelist download
  - OSPF for dynamic routing
- Managed devices terminating VIA connection
- Palo Alto Network Integration
- IP address assignment to managed devices using DHCP and VLAN pool configuration
- DHCP relay using IPv6 helper address.
- The following WAN Uplink features:
  - Dynamic Path Selection
  - IP Health Check
  - Uplink load balancing
  - WAN optimization
  - Cellular
- Communication with server over the following standard protocols:
  - Kerberos
  - NTLM
  - WISPr
  - LDAP
  - Dynamic DNS
- EST for certificate enrollment

Enabling IPv6

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

The following procedure describes how to enable the IPv6 option:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > Firewall** tab.
2. Expand the **Global Settings** accordion.
3. Click **IPv6 enable** check box to enable the IPv6 option.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Enabling IPv6 Support for Mobility Master and APs

This release of ArubaOS provides IPv6 support for a Mobility Master and access points. You can now configure the Mobility Master with an IPv6 address to manage the managed devices and APs. Both IPv4 and IPv6 APs can terminate on the IPv6 managed device. You can provision an IPv4 or IPv6 AP in the network when the managed device interface is configured with an IPv6 address. An IPv6 AP can serve both IPv4 and IPv6 clients.

Starting from ArubaOS 8.4.0.0, you can configure either or both IPv4 and IPv6 addresses in one data zone of an AP MultiZone profile. For more information on configuring AP MultiZone profile, see Configuring MultiZone on page 456.



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

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

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

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

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

- `show ap debug`

The following table lists IPv6 features:

Table 22: *IPv6 APs Support Matrix*

Features	Supported on IPv6 APs
Forward Mode - Tunnel	Yes
Forward Mode - Decrypt Tunnel	Yes
Forward Mode - Bridge	Yes
Forward Mode - Split Tunnel	Yes
AP Type - Campus AP	Yes
AP Type - Remote AP	Yes
AP Type - Mesh Node	No
CPsec	Yes
Wired-AP or Secure-Jack	Yes
Fragmentation or Reassembly	Yes
MTU Discovery	Yes
Provisioning Through Static IPv6 Addresses	Yes
Provisioning Through IPv6 FQDN Master Name	Yes
Provisioning from WebUI	Yes
Provisioning Through DHCPv6 Option 52	Yes
AP Boot by Flash	Yes
AP Boot by TFTP	No
WMM QoS	No
AP Debug and Syslog	Yes
ARM & AM	Yes
WIDS	Yes (Limited)
CLI Support for Users and Datapath	Yes
AP MultiZone Profile	Yes
Duplicate Address Detection (DAD)	Yes

Configuring IPv6 Addresses

You can configure IPv6 addresses for the management interface, VLAN interface, and the loopback interface of the Mobility Master and managed device. Up to three IPv6 addresses can be configured for each VLAN interface. The IPv6 address configured on the loopback interface or the first VLAN interface becomes the default IPv6 address of the device.



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

As per Internet Assigned Numbers Authority (IANA), a managed device supports IPv6 addresses beyond the following ranges:

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

The following procedure describes how to configure link local address.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces** > **VLANs** tab.
2. Select a Vlan from the **VLANs** table.
3. Select the corresponding Vlan Id from the **VLANS <Vlan name>** table.
4. Select the **IPv6** tab and expand the **IP Address Assignment** accordion.
5. Enable **Use Static addresses. IPV6 addresses** table appears.
6. Click **+** in IPV6 addresses.
7. In the **Add New IPV6 addresses** pop-up, select **Link-local** for **Address Type.** Enter the **IPV6 address** and select the **Use EUI-64 Format** check box, if applicable. Click **OK**.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following procedure describes how to configure global unicast address.

1. Select a VLAN from the **VLANs** table.
2. Select the corresponding VLAN ID from the **VLANS <Vlan name>** table.
3. Select the **IPV6** tab and expand the **IP Address Assignment** accordion.
4. Enable **Use Static addresses. IPV6 addresses** table appears.
5. Click **+** in IPv6 addresses.
6. In the **Add New IPV6 addresses** pop-up, select **Global unicast** for **Address Type.** Enter the **IPV6 address** and select the **Use EUI-64 Format** check box, if applicable. Click **OK**.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following procedure describes how to configure loopback interface address.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System** > **General** tab.
2. Select **Loopback Interface** accordion and enter the loopback address in the **IPv6 address** field.
3. Click **Submit**.

4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.



You cannot configure the management interface address using the WebUI.

The following CLI commands configure the link local address.

```
(host) [md] (config)#interface vlan <id>
(host) [md] (config-submode)#ipv6 address link-local <ipv6-address>
```

The following CLI commands configure the global unicast address.

```
(host) [md] (config)#interface vlan <id>
(host) [md] (config-submode)#ipv6 address <ipv6-prefix>/<prefix-length>
```

The following CLI commands configure configure the global unicast address (EUI 64 format).

```
(host) [md] (config)#interface vlan <id>
(host) [md] (config-submode)#ipv6 address <ipv6-prefix/prefix-length> eui-64
```

The following CLI commands configure the management interface address.

```
(host) [md] (config)#interface mgmt
(host) [md] (config-submode)#ipv6 address <ipv6-prefix/prefix-length>
```

The following CLI commands configure the loopback interface address.

```
(host) [md] (config)#interface loopback
(host) [md] (config-submode)#ipv6 address <ipv6-prefix>
```

Configuring IPv6 Static Neighbors

The following procedure describes how to configure IPv6 static neighbors.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces** > **IPv6 Neighbors** tab.
2. Click **+** in the **IPv6 Neighbors** table and enter the following details:
3. IPv6 address
4. Link-layer addr
5. Select a **VLAN interface** from the drop-down list.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures a static neighbor on a VLAN interface.

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

Configuring IPv6 Default Gateway and Static IPv6 Routes

The following procedure describes how to configure IPv6 default gateway.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces > IP Routes** tab.
2. Click + under the **Static Default Gateway** accordion.

3. Select **IPv6** from the **IP version** drop-down list, and enter the IPv6 address in the **IP address** field.
4. Click **Submit** to add the address to the IPv6 default gateway table.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following procedure describes how to configure static IPv6 routes.

1. Expand the **IP Routes** accordion.
2. Click + under the **IP Routes** table.
3. Select **IPv6** from the **IP version** drop-down list.
4. Enter the **Destination IP address** and the **Forwarding settings**.
5. Click **Submit** to add the static route to the IPv6 routes table.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures the IPv6 default gateway.

```
(host) [md] (config)#ipv6 default-gateway <ipv6-address> <cost>
```

The following CLI command configures static IPv6 routes.

```
(host) [md] (config)#ipv6 route <ipv6-prefix/prefix-length> <ipv6-next-hop> <cost>
<ipv6-next-hop>  = X:X:X:X::X
```

Configuring Prefix Delegation

Prefix delegation can be used to assign a network address prefix to a customer site, as defined in IPv6 prefix delegation protocol (RFC 3769). The hosts at the customer site use this prefix to derive a unique IPv6 address using RA and SLAAC. Prefix delegation client uses DHCPv6 IA_PD to request and assign prefixes.

As part of addition of prefix delegation, the following features are supported on the managed device:

- IPv6 address can be assigned using DHCPv6 IA_NA. For DHCPv6 stateful address assignment, DHCP client process is started for the interface VLAN to retrieve and manage the address.
- PD client is supported to retrieve the prefix from ISP using DHCPv6 IA_PD, PD client process is started for interface VLAN to retrieve and manage the IA_PD lease.
- A PD-based address, based on the prefix obtained on uplink VLAN using PD client can be configured on other interface VLANs. When a PD-based address is configured, that prefix is advertised using RA on that VLAN. This RA helps the host to derive a unique SLAAC address from the prefix advertised.



PD client and DHCPv6 client are not allowed to be configured in different VLANs. PD-based address cannot be configured on an interface that is configured to run PD or DHCPv6 clients.

Execute the following commands in the CLI to automate prefix delegation, and stateful IPv6 address configuration:

IPv6 address configuration under interface vlan

```
(host) [md] (config)#interface vlan 101
(host) [md] (config-submode)#ipv6 address dhcp6-client
            pd X:X:X:X::X/<0-128>
(host) [md] (config-submode)#ipv6 dhcp pdclient <pd_name>
```

Prefix delegation configuration under interface vlan

```
(host) [md] (config-submode)#ipv6 dhcp
           pdclient
           server
(host) [md] (config-submode)#ipv6 dhcp pdclient
            <pd_name>
```

IPv6 PD-based address configuration under downstream vlan interface

```
(host) [md] (config)#interface vlan 101
(host) [md] (config-submode)#ipv6 address
           dhcp6-client
           link-local
           pd X:X:X:X::X/<0-128>
(host) [md] (config-submode)#ipv6 address pd <pd_name> ::X:X:X:X:X
```

IPv6 PD status

```
(host)[md] #show ipv6 pd status
DHCPv6 PD Client is enabled
Uplink VLAN     : 100
Label           : site1
Prefix          : 2001:0:3::/48
65536 unique /64 prefixes are derivable from the acquired IA PD lease
Preferred lifetime 604800s, Valid lifetime 2592000s
Last request/renewal for the lease done at Thu Apr 14 04:46:15 2016
Lease expires at Sat May 14 04:46:15 2016
Downlink VLANs
--------------
VlanId  Prefix
------  ------
101      2001:0:3:12:1:2:3:4/64
```

Managing IP Addresses

You can change the default managed device IP address by assigning a different VLAN interface address or the loop back interface address. You can also turn on Syslog messaging for IPv6 (similar to IPv4 logging) using the `logging <ipv6 address>` command. For more information on logging, see Configuring Logging on page 986.

The following procedure describes how to configure the default managed device IP address:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System** > **General** tab.
2. Expand the **Controller IP address** accordion.
3. Select the VLAN Id or the loopback interface ID from the **IPv6 address** drop-down list.



In a native IPv6 deployment scenario, the configuration of IPv4 address is optional. Hence, you can select **None** from the **IPv4 address** drop-down list to remove the IPv4 address of the Mobility Master, when IPv6 address is configured on the Mobility Master.

4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure an IPv6 address to the managed device:

```
(host) [md] (config)#controller-ipv6 loopback
(host) [md] (config)#controller-ipv6 vlan <id> address <ipv6 address>
```

The following CLI command enables logging over IPv6

```
(host) [md] (config)#logging <ipv6 address>
```

The following CLI command removes the IPv4 address of the Mobility Master, when a valid IPv6 address is configured:

```
(host) [mynode] (config) #no controller-ip
```

Configuring Multicast Listener Discovery

You can enable the IPv6 multicast snooping on the managed device by using the WebUI or CLI and configure MLD parameters such as query interval, query response interval, robustness variable, and ssm-range.

The Source Specific Multicast (SSM) supports delivery of multicast packets that originate only from a specific source address requested by the receiver. You can forward multicast streams to the clients if the source and group match the client subscribed source group pairs (S,G).

The managed device supports the following IPv6 multicast source filtering modes:

- Include - In Include mode, the reception of packets sent to a specified multicast address is enabled only from the source addresses listed in the source list. The default IPv6 SSM address range is FF3X::4000:1 – FF3X::FFFF:FFFF, and the hosts subscribing to SSM groups can only be in the Include mode.
- Exclude - In Exclude mode, the reception of packets sent to a specific multicast address is enabled from all source addresses. If there is a client in the Exclude mode, the subscription is treated as an MLDv1 join.

For more information on MLD feature, see **RFC 3810** and **RFC 4604**.

MLD snooping does not add IPv6 Solicited-Node multicast address or groups to the multicast table. A Solicited-Node multicast address is an IPv6 multicast address valid within the local-link (example, an Ethernet segment or a Frame Relay cloud). Every IPv6 host has at least one such address per interface. Solicited-Node multicast addresses are used in Neighbor Discovery Protocol for obtaining the layer 2 link-layer addresses of other nodes.

The following procedure describes how to modify IPv6 MLD snooping.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces >VLANs** tab.
2. Click a VLAN name under **VLANs** and click the corresponding VLAN ID under **Vlan Ids**.
3. Click **IPv6** tab.
4. Expand the **Multicast Listener Discovery (MLD)** accordion.
5. Select **snooping** from the **MLD** drop-down list to enable IPv6 MLD snooping.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following procedure describes how to modify IPv6 MLD parameters.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces** > **Multicast** tab.
2. Expand the **MLD** accordion.
3. Enter the required values in the following fields:
4. **Robustness variable**: default value is 2

5. **Query interval**: default value is 125 seconds
6. **Query response interval**: default value is 100 (1 or 10 seconds).
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following procedure describes how to configure the SSM range.

1. In the **Managed Network** node hierarchy, navigate to **Configuration** > **Interfaces > Multicast** tab.
2. Expand the **MLD** accordion.
3. Use the **SSM range start-ip** and **SM range mask-ip** fields to configure the SSM Range.
4. Click **Submit** to save your changes.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

To enable IPv6 MLD snooping:

```
(host) [md] (config) #interface vlan 1
(host) [md] (config-submode)#ipv6 mld snooping
```

To view if IPv6 MLD snooping is enabled:

```
(host) [md] (config-submode)#show ipv6 mld interface
```

To view the MLD Group information:

```
(host) [md] (config)#show ipv6 mld group
```

To modify IPv6 MLD parameters:

```
(host) [md] (config)#ipv6 mld
(host) [md] (config-mld) # query-interval <time in seconds (1-65535)>|query-response-
interval <time in 1/10th of seconds (1-65535)|robustness-variable <value (2-10)>
```

To view MLD configuration:

```
(host) [md](config-submode)#show ipv6 mld config
```



When you enter the SSM range ensure that the upstream router has the same range, else the multicast stream would be dropped.

Dynamic Multicast Optimization

When multiple clients are associated to an AP and when one client is subscribed for a multicast stream, all the clients associated to the AP receive the stream, as the packets are directed to the multicast MAC address. To restrict the multicast stream to only the subscribed clients, DMO sends the stream to the unicast MAC address of the subscribed clients. DMO is currently supported for both IPv4 and IPv6.

The following CLI commands configure DMO.

```
(host) [md] (config) #wlan virtual-ap default
(host) [md] (Virtual AP profile "default") #dynamic-mcast-optimization
```

The following CLI command verifies the DMO configuration.

```
(host) #show wlan virtual-ap
```

Decrypt-Tunnel DMO

Starting from ArubaOS 8.2.0.0, Decrypt-Tunnel DMO support for IPv6 clients is introduced. This enhancement ensures that ArubaOS optimizes the multicast traffic for IPv6 wireless clients in the D-Tunnel mode. DDMO is done by converting multicast streams to unicast streams in the AP rather than in the controller or managed device.

A sample use case is when a wired multicast server sends an IPv6 multicast stream to IPv6 wireless clients that are associated to a Decrypt-Tunnel Virtual AP on an AP.

Important Points to Note

- DDMO is supported for both IPv4 and IPv6.
- If the number of wireless clients subscribed for the multicast traffic is more than the DDMO threshold, multicast traffic is not converted to unicast traffic; instead, the multicast traffic is sent to air directly.
- Because the conversion takes place in the AP instead of the controller or the managed device, make sure that the multicast packets can go through AP datapath.



Because the conversion takes place in the AP instead of the controller or the managed device, make sure that the multicast packets can go through AP datapath.

Troubleshooting

Execute any of the following commands for troubleshooting DDMO-related issues:

- `show ap debug aid-table bssid <bssid> advanced`—shows the association ID table
- `show ipv6 mld proxy-group`— displays MLD proxy-group details
- `show ipv6 mld group`—displays MLD group details
- `show datapath ipv6-mcast group`—displays the IPv6 multicast group
- `show datapath ipv6-mcast station`—displays the IPv6 station membership
- `show datapath ipv6-mcast destination`—displays the IPv6 tunnel and port membership

Limitations

The following are the MLDv2 limitations:

- Managed Device cannot route multicast packets.
- For mobility clients mld proxy should be used.
- VLAN pool scenario stream is forwarded to clients in both the VLANs even if the client from one of the VLANs is subscribed.
- DMO is applicable for wired clients in managed device.

Debugging IPv6

ArubaOS provides the following debug commands for IPv6:

- `show ipv6 global` — displays if IPv6 is enabled globally or not
- `show ipv6 interface` — displays the configured IPv6 address, and any duplicate addresses
- `show ipv6 route` or `show datapath route ipv6` — displays the IPv6 routing information
- `show ipv6 ra status` — displays the RA status

- `show Datapath session ipv6` — displays the IPv6 sessions created, and the sessions that are allowed
- `show datapath frame` — displays the IPv6 specific counters

The following procedure describes how to use the ping and tracepath options.

1. To ping an IPv6 host, in the **Mobility Master** node hierarchy**,** navigate to **Diagnostics > Tools > Ping** tab, enter an IPv6 address, and click **Ping**.
2. To trace the path of an IPv6 host, in the **Mobility Master** node hierarchy**,** navigate to **Diagnostics > Tools > Tracepath** tab, enter an IPv6 address, and click **Trace**.

The following CLI commands ping an IPv6 host.

```
(host)#ping ipv6 <global-ipv6-address>
(host)#ping ipv6 interface vlan <vlan-id> <linklocal-address>
```

The following CLI command traces the path of an IPv6 host.

```
(host)#tracepath <global-ipv6-address>
```

Provisioning an IPv6 AP

You can provision an IPv6 AP on an IPv6 Mobility Master. You can either configure a static IP address or obtain a dynamic IPv6 address via stateless-autoconfig. The managed device can act as the default gateway for the IPv6 clients, if static IPv6 routes are set on the managed device.



A wired client can now connect to the Ethernet interface of an IPv6 enabled AP.

Execute the following command to provision a static IPv6 address:

```
(host) [mynode](config)# provision-ap
```

Enhancements to IPv6 Support on AP

ArubaOS provides the following IPv6 enhancements on the AP:

- DNS based IPv6 master discovery
- FTP support for image upgrade in an IPv6 network
- DHCPv6 client support
- Bridge mode IPv6 firewall support

Filtering an IPv6 Extension Header

ArubaOS firewall is enhanced to process the IPv6 Extension Header to enable IPv6 packet filtering. You can now filter the incoming IPv6 packets based on the EH type. You can edit the packet filter options in the default EH, using the CLI. The default EH alias permits all EH types.

The following CLI commands permit or deny IPv6 packets matching an EH type.

```
(host) [md](config) #netexthdr default
(host) [md](config-exthdr) #eh <eh-type> permit | deny
```

The following CLI command shows the EH types denied.

```
(host) [md](config-exthdr) #show netexthdr default
```

Configuring a Captive Portal over IPv6

IPv6 is now enabled on the captive portal for user authentication on the Aruba managed device. For user authentication, use the internal captive portal that is initiated from the managed device. A new parameter **captive** has been added to the IPv6 captive portal session ACL:

```
(host) [md] (config)#ipv6 user alias controller 6 svc-https captive
```



Captive portal authentication, customization of pages, and other attributes are same as IPv4.

You can configure captive portal over IPv6 (similar to IPv4) using the WebUI or CLI. For more information on configuration, see Configuring Captive Portal in the Base Operating System on page 329.

Working with IPv6 RAs

ArubaOS enables the managed device to send RA in an IPv6 network. Each host auto generates a link local address when you enable ipv6 on the host. The link local address allows the host to communicate between the nodes attached to the same link.

The IPv6 stateless autoconfiguration mechanism allows the host to generate its own addresses using a combination of locally available information and information advertised by the routers. The host sends a router solicitation multicast request for its configuration parameters in the IPv6 network. The source address of the router solicitation request can be an IP address assigned to the sending interface, or an unspecified address if no address is assigned to the sending interface.

The routers in the network respond with an RA. The RAs can also be sent at periodic intervals. The RA contains the network part of the Layer 3 IPv6 address (IPv6 Prefix). The host uses the IPv6 prefix provided by the RA; it generates the universally unique host part of the address (interface identifier), and combines the two to derive the complete address. To establish continuous connectivity to the default router, the host starts the neighbor reachability state machine for the router.



ArubaOS uses Radvd, an open source Linux IPv6 RA daemon maintained by Litech Systems Design.

ArubaOS now provides support for DNS Search List (DNS-SL) option through IPv6 Router Advertisements that allows the IPv6 clients to resolve incomplete domain names. For example, if the DNS-SL is *template.com* and the incomplete domain name is *example*, then the client application resolves the issue automatically by appending DNS-SL to the incomplete domain name. Hence, the final domain name is modified to *example.template.com*.

You can perform the following tasks on the managed device to enable, configure, and view the IPv6 RA status on a VLAN interface:

- Configure IPv6 RA on a VLAN
- Configure Optional Parameters for RA
- Configure neighbor discovery reachable time
- Configure neighbor discovery retransmit time
- Configure RA DNS
- Configure RA DNS-SL
- Configure RA hop-limit
- Configure RA interval
- Configure RA lifetime
- Configure RA managed configuration flag

- Configure RA MTU
- Configure RA other configuration flag
- Configure RA Preference
- Configure RA prefix
- View IPv6 RA Status

Configuring an IPv6 RA on a VLAN

You must configure the IPv6 RA functionality on a VLAN for it to send solicited or unsolicited RAs on the IPv6 network. You must configure the following for the IPv6 RA to be operational on a VLAN:

- IPv6 global unicast address
- Enable IPv6 RA
- IPv6 RA prefix

Important Points to Note

- The advertised IPv6 prefix length must be 64 bits for the stateless address autoconfiguration to be operational.
- You can configure up to three IPv6 prefixes per VLAN interface.
- Each IPv6 prefix must have an on-link interface address configured on the VLAN.
- Ensure you configure the upstream routers to route the packets back to Aruba managed device.

The following procedure describes how to configure the IPv6 RA on a VLAN:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces > VLANs** tab.
2. Select a VLAN name under the **VLANs** table.
3. Select the corresponding VLAN ID from the **VLANs <name>** table and click the **IPv6** tab.
4. To configure an IPv6 global unicast address, follow the steps below:
   a. Expand the **IP Address Assignment** accordion.
   b. Enable **Use Static addresses. IPV6 addresses** table appears.
   c. Click **+** in IPV6 addresses.
   d. In the **Add New IPV6 addresses** pop-up, select **Global unicast** for **Address Type.** Enter the **IPV6 address** and select the **Use EUI-64 Format** check box, if applicable. Click **OK**.
5. To enable an IPv6 RA on a VLAN, select the **Router advertisements (ra)** check box under **Neighbor Discovery** accordion.
6. To configure an IPv6 RA prefix for a VLAN, follow the steps below:
   a. Click **+** in the **RA prefixes** table in the **Neighbor Discovery** accordion.
   b. Enter a value in the **IPv6 RA** field.
   c. Click **OK**.

   You can add up to three IPv6 prefixes per VLAN interface.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure RA on a VLAN:

```
(host) [md](config) #interface vlan <vlanid>
(host) [md](config-subif)#ipv6 address <prefix>/<prefix-length>
(host) [md](config-subif)#ipv6 nd ra enable
(host) [md](config-subif)#ipv6 nd ra prefix X:X:X:X::X/64
```

Configuring Optional Parameters for RAs

In addition to enabling the RA functionality, you can configure the following IPv6 neighbor discovery and RA options on a VLAN:

- Neighbor discovery reachable time – The time, in milliseconds, that a node assumes a neighbor is reachable after receiving a reachability confirmation.
- Neighbor discovery retransmit time – The time, in milliseconds, between retransmitted Neighbor Solicitation messages.
- RA DNS – The IPv6 recursive DNS Server for the VLAN.



On Linux systems, clients must run the open rdnssd daemon to support the DNS server option. Windows 7 does not support the DNS server option.

- RA DNS-SL – The IPv6 recursive DNS Server Search List for the VLAN.



You can configure up to 8 DNS-SLs for each VLAN, and each DNS-SL supports a maximum of 32 characters.

- RA hop-limit – The IPv6 RA hop-limit value. It is the default value to be placed in the Hop Count field of the IP header for outgoing (unicast) IP packets.
- RA interval – The maximum and minimum time interval between sending unsolicited multicast RA from the interface, in seconds.
- RA lifetime – The lifetime associated with the default router in seconds. A value of zero indicates that the router is not a default router and will not appear on the default router list. The router lifetime applies only to the router's usefulness as a default router; it does not apply to information contained in other message fields or options.
- RA managed configuration flag (Enable DHCP for address) – A flag that indicates that the hosts can use the DHCP server for address autoconfiguration besides using RAs.
- RA MTU – The MTU that all the nodes on a link use.
- RA other configuration flag (Enable DHCP for other information – A flag that indicates that the hosts can use the administered (stateful) protocol for autoconfiguration of other (non-address) information.
- RA preference – The preference associated with the default router.

You can use the WebUI or CLI to configure these options.



It is recommended that you retain the default value of the RA interval to achieve better performance.

If you enable RAs on more than 100 VLAN interfaces, some of the interfaces may not send out the RAs at regular intervals.

The following procedure describes how to configure the optional parameters:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces >VLANs** tab.
2. Select a VLAN name under the **VLANs** tab.
3. Select the corresponding VLAN ID from the **VLANs <name>** table and click **IPv6** tab.
4. Expand the **Neighbor Discovery** accordion and configure the following neighbor discovery and RA options for the VLAN based on your requirements:
   a. Select the **RA advertisements(ra)** check box to enable RA advertisements.
   b. Enter a value in the **Reachable time (ms)** field. The allowed range is 0-3,600,000 msec. The default value is zero.
   c. Enter a value in the **Retransmit time(ms)** field. The allowed range is 0-3,600,000 msec. The default value is zero.
   d. Enter a hop-limit value in the **RA hop limit** field. The allowed range is 1-255. The default value is 64.
   e. Enter a value in the **RA interval minimum (sec) and RA interval maximum (sec)** fields. Allowed range is 3-0.75 times the maximum RA interval value in seconds. The default minimum value is 0.33 times the maximum RA interval value
   f. Enter a value in the **RA lifetime** (sec) field. A value of zero indicates that the router is not a default router. Apart from a zero value, the allowed range for the lifetime value is the RA interval time to 9,000 seconds. The default and minimum value is three times the RA interval time.
   g. Click **+** in **RA Prefixes**, enter **IPV6 RA** and **IPV6 RA Prefix**. Click **OK**.
   h. Click **+** in the **Recursive DNS servers** field, and enter a **DNS IPV6 address** in the **Add DNS Server** pop-up. Click **OK**.
   i. Select **Enabled** from the **DHCP for address** drop-down list to enable the hosts to use the DHCP server for address autoconfiguration apart from any addresses auto configured using the RA.
   j. Enter a value in the **RA MTU option** field. The allowed range is 1,280-maximum MTU allowed for the link.
   k. Click **DHCP for other info** check box to enable the hosts to use the DHCP server for autoconfiguration of other (non-address) information.
   l. Select the **Router preference**as **High**, **Medium**, or **Low**.
   m. Click **+** in the **DNS Search List** field, and enter a **Domain Name** in the **Add a DNS Suffix Domain Name** pop-up. Click **OK**.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Execute the following CLI commands to configure the neighbor discovery and RA options for a VLAN interface:

To configure neighbor discovery reachable time:

```
(host) [md] (config) #interface vlan <vlan-id>
(host) [md] (config-submode)#ipv6 nd reachable-time <value>
```

To configure neighbor discovery retransmit time:

```
(host) [md] (config-submode)#ipv6 nd retransmit-time <value>
```

To configure IPv6 recursive DNS server:

```
(host) [md] (config-submode)#ipv6 nd ra dns X:X:X:X::X
```

To configure RA hop-limit:

```
(host) [md](config-submode)#ipv6 nd ra hop-limit <value>
```

To configure RA interval:

```
(host) [md] (config-submode)#ipv6 nd ra interval <value> <min-value>
```

To configure RA lifetime:

```
(host) [md] (config-submode)#ipv6 nd ra life-time <value>
```

To enable hosts to use DHCP server for stateful address autoconfiguration:

```
(host) [md] (config-submode)#ipv6 nd ra managed-config-flag
```

To configure MTU for RA:

```
(host) [md] (config-submode)#ipv6 nd ra mtu <value>
```

To enable hosts to use DHCP server for other non-address stateful autoconfiguration:

```
(host) [md] (config-submode)#ipv6 nd ra other-config-flag
```

To specify a router preference:

```
(host) [md] (config-submode)#ipv6 nd ra preference [High | Low | Medium]
```

To view the IPv6 RA status on the VLAN interfaces:

```
(host) [md]#show ipv6 ra status
```

To configure the DNS-SL on the VLAN interface:

```
(host) [md] (config) #interface vlan <vlan-id>
(host) [md] (config-submode)#ipv6 nd ra dns-sl <value>
```

IPv6 RA Proxy

An IPv6 network deployment typically has one or more upstream routers to delegate IPv6 prefixes through RA to clients. When a client connects to the network, it would begin with sending Router Solicitations or DHCP IPv6 requests. In case of SLAAC using RA where clients sends Router Solicitations, upstream routers can either respond with unicast (L2 or L3) RA or with multicast RA. Whenever a new client joins the network, a unicast or a multicast RA is sent to from the router to the client. If it is a multicast packet then existing clients also receive the RA, which results in increasing the traffic. Starting from ArubaOS 8.1, this issue is addressed by enabling IPv6 proxy RA to snoop incoming unsolicited RA and Router Solicitations packets.

The following procedure describes how to configure IPv6 RA proxy settings:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > Firewall** tab.
   a. Expand the **Global Settings** accordion.
2. Select **IPV6 Proxy RA** check box.
3. Click **Submit**.

4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Execute the following command to configure an interval for proxy RA:p

```
(host) [md] (config) #ipv6 proxy-ra interval
```

Execute the following command to enable the proxy RA:

```
(host) [md] (config) #ipv6 proxy-ra
```

Execute the following command to view the status of the IPv6 proxy RA:

```
(host) #show ipv6 ra proxy
IPv6 RA Proxy status: enabled
IPv6 RA Proxy interval: 600
```

Centralized Licensing Support for IPv6

ArubaOS now supports centralized licensing architecture in IPv6 network also, where a local controller containing IPv4 or IPv6 address acts as the license client and communicates with the license server containing IPv6 address to obtain the available licenses. The centralized licensing information is sent between license server and license client through heartbeat messages based on UDP. With the introduction of IPv6 address support, the heartbeat messages between the license server and license client must use IPv6 address as source and destination IP address.



The license server now supports both IPv4 and IPv6 address clients. However, you cannot configure both IPv4 and IPv6 addresses and must configure either IPv4 address or IPv6 address as license server IP address.

Associating Mobility Master or Stand-alone Controller to External License Server

You can associate one or more stand-alone controllers or Mobility Masters to an external license server, allowing the controller or Mobility Master to obtain licenses from a license pool on that external server. If Mobility Master connects to a license server, that Mobility Master then acts a proxy license server, distributing licenses to the managed devices that are part of the root licensing pool on the Mobility Master server.

If you use an external license server, all primary and backup Mobility Master servers, and standalone controllers must be able to communicate with the external license server. Managed devices associated to Mobility Master do not need connectivity with the license server.

You can connect one or more stand-alone controllers or Mobility Masters to an external license server in the following scenarios:

- Single Mobility Master and Local Controller in IPv6 Network
- Single Mobility Master and Local Controller in Mixed Network
- Multiple Mobility Masters in IPv6 network
- Multiple Mobility Masters in Mixed Network

Single Mobility Master and Local Controller in IPv6 Network

The centralized licensing feature is supported in a single Mobility Master that acts as a centralized license server configured with IPv6 controller IP address. A local controller works as a license client that is configured with IPv6 controller IP address as shown in the following table:

Table 23: *Mobility Master and Local Controller in IPv6 matrix*

License Server Controller IP	Licenses from Mobility Master	Compatible	License heartbeat transport from client to master
IPv4 + IPv6 address	IPv4 + IPv6 address	Yes	IPv6 address

Single Mobility Master and Local Controller in Mixed Network

The **CFGM** process of the License Manager upgrades the Mobility Master license server from IPv4 to IPv6 address before the Mobility Master license client IP address is upgraded to IPv6 address. The license server can have both IPv4 and IPv6 controller IP addresses while the license client can have IPv4 controller IP address only as shown in the following table:

Table 24: *Mobility Master and Local Controller in IPv4 IPv6 matrix*

License Server Controller IP	Licenses from Mobility Master	Compatible	License heartbeat transport from client to master
IPv4 address only	IPv4 address only	Yes	IPv4 address
IPv4 + IPv6 address	IPv4 address only	Yes	IPv4 address



The License Manager supports IPv6 address only when you configure IPv4 address of the managed device. If no license server IP address is configured, the license client uses either IPv4 or IPv6 address of license server based on master IPv4 or master IPv6 address.

Multiple Mobility Masters in IPv6 network

The centralized licensing feature is now supported for multiple Mobility Masters where a single Mobility Master acts as a license server that is configured with the IPv6 address of the controller. The remaining Mobility Masters act as relay servers that support configuration of IPv6 license server IP addresses pointing to the IPv6 license server as shown in the following table:

Table 25: *Multiple Mobility Masters in IPv6 matrix*

License Server Controller IP	License Relay Server License Server IP	License Relay Server Controller IP	Compatible	License heartbeat transportfrom license relay server to license server
IPv4 + IPv6 address	IPv6 address only	IPv4 + IPv6 address	Yes	IPv6 address



The relay servers and their license client controllers follow single Mobility Master and local IPv6 support network.

Multiple Mobility Masters in Mixed Network

The centralized licensing feature for multiple Mobility Masters in a mixed network is supported in the following scenarios:

- When the license server is configured with IPv4 address of controller, the remaining Mobility Masters acting as relay servers can only work with IPv4 license server. However, the relay server itself can have IPv4 or IPv6 controller IP configured.
- When the license server is configured with IPv6 address of controller, the remaining Mobility Masters acting as relay servers can work with IPv4 license server; and the relay server itself can have IPv4 or a combination of IPv4 and IPv6 controller IP configured.

The following table describes the scenarios:

Table 26: *Multiple Mobility Masters in IPv4 IPv6 matrix*

License Server Controller IP	License Relay Server License Server IP	License Relay Server Controller IP	Compatible	License heartbeat transport from license relay server to license server
IPv4 address only	IPv4 address only	IPv4 + IPv6 address	Yes	IPv4 address
IPv4 + IPv6 address	IPv4 address only	IPv4 + IPv6 address	Yes	IPv4 address

The following procedure describes how to associate a Mobility Master or stand-alone controller to an external license server.

1. Before you begin, access the command-line interface Mobility Master and remove any unnecessary license pool profiles.
2. From the Mobility Controller configuration node on a standalone controller, or from the Mobility Master configuration node for a Mobility Master, navigate to **Configuration > License** page.
3. In the **License Management** option, select **External license server** radio button.
4. The **External License Server** fields are displayed.
5. Select either **IPv4** or **IPv6** radio button based on your preference.
6. In the **IP address** field, enter the IPv4 or IPv6 address of the external license server.
7. Click **Submit** to save your changes.

To associate a Mobility Master with an external license server, run the **no license-pool-profile <profile>** command from the Mobility Master (mm) configuration node to remove any local license pools. Then, run the **license server-ip<ip-addr>|<ipv6-addr>** command to define the external server.

For example:

```
(host)[mm] (config) #no license-pool-profile /USA/southwest
(host)[mm] (config) #no license-pool-profile /USA/northeast
(host)[mm] (config) #no license-pool-profile /APAC/India
(host)[mm] (config) #no license-pool-profile /USA/Beijing
(host)[mm] (config) #license server-ip  2002::6
```

To view license usage details for standalone controllers or Mobility Master servers connected to an external licensing pool, run the **show license-usage client** command.

To associate a standalone controller to an external license server, run the **license server-ip<ip-addr>|<ipv6-addr>** command to define the external server.

For example:

```
(host)[mm] (config) #license server-ip  10.1.1.91
```

IPsec Support

IPsec support is enhanced to accommodate IPv6 which includes overlay networks across IPv4 and IPv6 IPsec Tunnels. IPsec is the base for security features like Site-to-Site VPNs, CPsec, Remote AP, and Master-Local deployments. The control plane handles the configuration of these features and translation to IPsec SA. The data plane handles encryption or decryption, encapsulation or decapsulation, tunneling, session setup, management, and routing of IPsec data.

Prior to this release only Global Unicast, Unique Local Unicast, and Link Local subnets were supported. Starting from ArubaOS 8.2.0.0, there is no limitation on the range that can configured for IPv6 IPsec Site-to-Site, and this enhancement has been extended to all platforms.

Starting from ArubaOS 8.1.0.0, IKEv2 or IPsec support is extended to IPv6 for the following topologies:

- Mobility Master-Managed Devices on page 143
- Control Plane Security (Tunnel Mode and D-Tunnel Mode) on page 147
- Remote AP (Tunnel Mode and D-Tunnel Mode) on page 148
- Remote AP(Split-Tunnel Forwarding Mode) on page 149
- Site-to-Site Crypto Map (Tunnel Mode and Transport Mode) on page 149

Mobility Master-Managed Devices

This topology is used to secure all control plane traffic exchanged between Mobility Master and a managed device. This is specifically used by **CFGM** process to push configuration from Mobility Master to the managed device.

While configuring IPv6 local and master addresses, assign and use an IPv4 or IPv6 address to accommodate all the other applications. For example, License Manager listens on IPv4 or IPv6 address that gets accommodated by Mobility Master- Managed Devices IPv6 deployment. For information on centralized licensing for both IPv4 and IPv6 addresses, license types, usage, and license installation, see *ArubaOS Licensing Guide*.

Configuring the Mobility Master IPv6 Address

You can configure the Mobility Master IPv6 address used for PSK authentication on the managed device, using the WebUI and CLI.



The initial IPv6 masteripv6 configuration on local must be done using the CLI. Once configured, any change in masteripv6 can be done through WebUI or CLI. This change in masteripv6 should remain in the same IP address family. Change from masteripv6 to masterip (vice-versa) requires a write erase on the managed device.

The following procedure describes how to configure the master IPv6 address for PSK authentication.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Controllers** menu.
2. Select a controller from the **Controllers** table.
3. Select **PSK** from the **Authentication** drop-down list.
4. Enter the **Mobility master IPV4 address**, **Mobility master IPV6 address,** and **Passphrase**.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures the master IPv6 address for PSK authentication.

```
(host) [md] (config) #masteripv6 <masteripv6>
ipsec <key> [fqdn <local-fqdn>][interface|{vlan <id>}][masteripv4 <masteripv4>]
```



In a native IPv6 deployment, it is not mandatory to configure the IPv4 address of the Mobility Master or managed device, when IPv6 address is configured. Hence, the **Mobility master IPV4 address** field in the WebUI and the **masteripv4** parameter under **masteripv6** CLI command are optional in a pure IPv6 network.

You can configure the master IPv6 address used for certificate-based authentication on the managed device.

The following procedure describes how to configure the master IPv6 address for certificate-based authentication.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Controllers** menu.
2. Select a controller from the **Controllers** table.
3. Select **Certificate** from the **Authentication** drop-down list.
4. Select **Custom/Factory** from the **Certificate type** drop-down list.
5. Enter the **Mobility master IPV4 address** and **Mobility master IPV6 address**.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures the master IPv6 address for certificate-based authentication:

```
(host) [myd] (config) #masteripv6 <masteripv6>
ipsec-custom-cert master-mac-1-c <mac-1-c> ca-cert <ca> server-cert [fqdn-v <local-
fqdn-v>][interface-c {uplink-v <uplink-v> | vlan-c <id-c>] masteripv4 {fqdn-
v|interface-c|masteripv4|suite-b} siute-b {fqdn-v|interface-c|masteripv4|suite-b}
ipsec-factory-cert master-mac-1 <MAC> [fqdn-c <local-fqdn-c>][interface-v {uplink-v
<uplink-v> | vlan-c <id-c>] [master-mac-2 <MAC>] [masteripv4 <masteripv4>]
```



ArubaOS reboots the managed device when the primary IPv6 address is changed on the managed device. However, a change in the secondary IPv6 address does not require a reboot of the managed device.

Configuring IPv4 or IPv6 Address for PSK Authentication

On Mobility Master, you can configure the IPv4 or IPv6 address of the managed device, to be used for PSK authentication:

The following procedure describes how to configure the IPv6 address of the managed device for PSK authentication:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Controllers** menu.
2. Click + under **Local Controllers IPSec Keys** table.
3. Select **IPSec Key** from the **Authentication** drop-down list.
4. Enter the **Local controller IPv4**, **Local controller IPv6 address**, and the **IPSec key** in the **Add New IPsec Controller** table.
5. Retype the IPsec key.
6. Click **Submit**.

7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Execute the following command in the CLI to configure the IPv6 address of the managed device for PSK authentication:

```
(host)[mynode](config)#localipv6 <local-switch-ipv6> ipsec <pre-shared-key> localipv4
<local-switch-ipv4>
```

You can configure the certificate-based authentication on the managed device.



In a native IPv6 deployment, it is not mandatory to configure the IPv4 address of the Mobility Master or managed device, when IPv6 address is configured. Hence, the **Local controller IPv4** field in the WebUI and the **localipv4** parameter under **localipv6** CLI command are optional in a pure IPv6 network.

Configuring IPv4 or IPv6 Address for Certificate-Based Authentication

On Mobility Master, you can configure the IPv4 or IPv6 address of the managed device, to be used for certificate-based authentication:

The following procedure describes how to configure the IPv6 address of the managed devices for certificate-based authentication.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Controllers** menu.
2. Click **+** in the **Local Controllers IPSec Keys** table.
3. Select **Certificate** from the **Authentication** drop-down list.
4. Enter the **Mac address** of the managed device.
5. Select **Factory** from the **Certificate type** drop-down list.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure certificate-based authentication on the managed device:.

```
(host) [mynode](config) #local-custom-cert local-mac <MAC> ca-cert <ca-cert-name>
server-cert <server-cert-name>
```

Monitoring and Managing Master Local IPv6 Settings

Execute the following command in the CLI, on Mobility Master to check the ipv6 address of the managed device:

```
(host) [mynode] (config) #show localipv6
Local Switches configured by Local Switch IPv6
----------------------------------------------
Switch IPv6 address of the Local  Corres IPv4 address of the Local  Key
--------------------------------  --------------------------------  ---
2002::1                           1.1.1.1                           ********
Execute  the following command in the CLI to check the IKE SA:
(host) [mynode] #show crypto isakmp sa
ISAKMP SA Active Session Information
------------------------------------
Initiator IP     Responder IP   Flags       Start Time      Private IP
------------     ------------   -----     ---------------   ----------
ISAKMP SA Active Session Information
------------------------------------
```

```
Initiator IP     Responder IP   Flags       Start Time      Private IP
------------     ------------   -----     ---------------   ----------
2002::1          2002::3        r-v2-p    Dec  4 15:15:31     -
Flags: i = Initiator; r = Responder
m = Main Mode; a = Agressive Mode; v2 = IKEv2
p = Pre-shared key; c = Certificate/RSA Signature; e =  ECDSA Signature
x = XAuth Enabled; y = Mode-Config Enabled; E = EAP Enabled
3 = 3rd party AP; C = Campus AP; R = RAP;  Ru = Custom Certificate RAP;

I = IAP
V = VIA; S = VIA over TCP
Total ISAKMP SAs: 1
Execute  the following command in the CLI to check the IPsec SA:
(host) [mynode] #show crypto ipsec sa
IPSEC SA (V2) Active Session Information
------------------------------------
Initiator IP     Responder IP     SPI(IN/OUT)        Flags Start Time        Inner IP
------------     ------------     ----------------   ----- ---------------   --------
2002::1          2002::3          cdb5f100/d533c500  T2    Dec  4 14:53:42     -
Flags: T = Tunnel Mode; E = Transport Mode; U = UDP Encap
L = L2TP Tunnel; N = Nortel Client; C = Client; 2 = IKEv2
Total IPSEC SAs: 1
```

Execute the following command in the CLI to retrieve the statistics of communication between Mobility Master and the managed device:

```
(host) [mynode] #show master-local stats 2002::1
Missed -> HB Req from Local(s)
------------------------------
IP Address  HB Req      HB Resp     Cfg Terminate  Peer Reset  Total Missed
----------  ------      -------     -------------  ----------  ------------
32.2.0.0    35155       35155       0              0           0


Last Sent Missed  Last Synced/First Missed
----------------  ------------------------
0                 Pending
```

Execute the following command in the CLI to check the progress of the configuration update:

```
(host)[mynode] #show switches state [complete|incomplete|inprogress|required]

(host) [mynode] (config) #show switches state complete
All Switches
------------
IP Address  IPv6 Address  Name             Location          Type  Model
Version
----------  ------------  ----             --------          ----  -----         -----
--
1.1.1.1     2002::1       abhi_vmc_61.122  Building1.floor1  LC    VMC-TACTICAL
8.0.0.0-svcs-ctrl_0000

Status       Configuration State  Config Sync Time (sec)  Config ID
----------   -------------------  ----------------------  ---------
up           UPDATE SUCCESSFUL    0                       22
Total Switches:1

(host) [mynode] (config) #show switches state incomplete
All Switches
------------
IP Address  IPv6 Address  Name  Location  Type  Model  Version  Status  Configuration
State
----------  ------------  ----  --------  ----  -----  -------  ------  --------------
-----
```

```
Config Sync Time (sec)  Config ID
----------------------  ---------
Total Switches:0

(host) [mynode] (config) #show switches state inprogress
All Switches
------------
IP Address  IPv6 Address  Name  Location  Type  Model  Version  Status  Configuration
State
----------  ------------  ----  --------  ----  -----  -------  ------  --------------
-----

Config Sync Time (sec)  Config ID
----------------------  ---------
Total Switches:0

(host) [mynode] (config) #show switches state required
All Switches
------------
IP Address  IPv6 Address  Name  Location  Type  Model  Version  Status  Configuration
State
----------  ------------  ----  --------  ----  -----  -------  ------  --------------
-----

Config Sync Time (sec)  Config ID
----------------------  ---------
Total Switches:0
```

Control Plane Security (Tunnel Mode and D-Tunnel Mode)

CPsec feature enables communication between an IPsec enabled AP (in CPsec mode) and a managed device. The configuration traffic between the managed device and AP (in CPsec mode) is routed through an IPsec tunnel, whereas the client traffic served by the AP is communicated to the managed device in clear. Heartbeats go in a GRE tunnel, even though they are locally generated.

To enable a range of IPv6 addresses for an AP that can terminate on a managed device, see Enabling a Range of IPv6 Addresses on page 147.

Enabling a Range of IPv6 Addresses

An AP can terminate on a Managed device, if Auto Cert Provisioning is enabled in the CPsec profile or if a range of IPv6 addresses are enabled under CPsec profile, and if Auto Cert Provisioning is disabled.

The following procedure describes how to enable a range of IPv6 addresses:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > CPSec** tab.
2. Click **Control Plane Security** accordion.
3. Click the **Enable CPSEC** toggle switch.
4. Click the **Enable auto Cert Provisioning** toggle switch.
5. Click the **Only accept APs from specified ranges** toggle switch.
6. Click **+** in the **Address ranges for Auto Cert Provisioning** table.
7. Enter the **Start address (Ipv4/Ipv6)** and **End address (IPv4/IPv6)** in the **New Address Range** dialogue box.
8. Click **OK**.
9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Execute the following command in the CLI to enable a range of IPv6 addresses:

```
(host) [md] (config) #control-plane-security
(host) [md] (Control Plane Security Profile) #auto-cert-allowed-addrs <startv6>
<endv6>
```

Execute the following command in the CLI to accept cert-provisioning for all the IPv6 addresses:

```
(host) [md] (config) #control-plane-security
(host) [md] (Control Plane Security Profile) #auto-cert-allow-all
```

You can check if the CPsec is enabled by executing the following command:

```
(host) [md] (Control Plane Security Profile) #show control-plane-security
```

Remote AP (Tunnel Mode and D-Tunnel Mode)

An AP terminating with an IPv6 address can be provisioned as a Remote AP using certs only. IPv6 L2TP pool is provisioned to assign IPv6 inner-ip address to AP. When a configuration request is initiated, the AP requests for IPv6 inner IP as the peer switch-ip.



In this release, only certificate-based Remote AP in forward-mode and decrypt mode tunnel is supported.

For information on provisioning inner IP address to Remote AP, see Provisioning an Inner IPv6 Address to a Remote AP on page 148.

Provisioning an Inner IPv6 Address to a Remote AP

The following procedure describes how to provision an inner IPv6 address to a Remote AP, by configuring L2TP IPv6 Pool address range.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand the **General VPN** accordion.
3. Click **+** in the **Address Pools** table.
4. Enter the **Pool Name**, **Start address**, and **End address** in the **Add New Address Pool** table.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Execute the following command in the CLI to provision an inner IPv6 address to a Remote AP, by configuring L2TP IPv6 Pool address range:

```
(host) [mynode] (config) #ipv6 local pool <pool_name_v6> <pool_start_addressv6> <pool_
end_addressv6>
```

Execute the following command in the CLI to view the total number of IPs in each pool and the IPs assigned from each pool:

```
(host) [mynode] (config) #show vpdn l2tp local pool
```

Execute the following command in the CLI to view the pools configured:

```
(host) [mynode] (config) #show vpdn l2tp configuration
```

Execute the following command in the CLI to provision an inner IPv6 address to a Remote AP, by configuring remote-ipv6 address in the whitelist DB entry:

```
(host) [mynode] (config) #whitelist-db rap add mac-address <mac address> ap-group <ap_
group> remote-ipv6 <remote_ipv6>
```

Remote AP(Split-Tunnel Forwarding Mode)

The Remote AP supports IPv6 clients in split-tunnel forwarding mode in a VAP profile. The STA in split-tunnel forwarding mode receives IPv6 address through RA (Router Advertisement) or DHCPv6 from the corporate network, and the controller forwards the traffic to the destination. Hence, AP datapath forwards the corporate traffic to the managed device over the GRE tunnel (split tunnel), and the Remote AP forwards the remaining traffic locally by using stateful source NAT. The traffic from STA is split based on the configured user role and session ACL applied to the STA.

For information on enabling the split-tunnel forwarding mode under VAP profile, see Understanding Split Tunneling on page 873.



Since Remote APs do not work on an IPv6 cluster, split-tunnel forwarding mode is not supported in an IPv6 cluster deployment.

Site-to-Site Crypto Map (Tunnel Mode and Transport Mode)

A VPN consists of multiple remote peers transmitting private data securely to one another over an unsecured network, such as the Internet. Site-to-site VPNs use tunnels to encapsulate data packets within normal IP packets for forwarding over IP-based networks; encryption ensures that privacy and authentication to ensure integrity of data. Listed below are the types of deployments supported in this release:

- IPv6 networks over IPv6 IPsec tunnel
- IPv6 networks over IPv4 IPsec tunnel
- IPv4 networks over IPv6 IPsec tunnel
- IPv4 and IPv6 networks over IPv4 IPsec tunnel
- IPv4 and IPv6 networks over IPv6 IPsec tunnel
- Static IPv6 Route to IPsec crypto-map
- IP compression support for IPv6 inner traffic



All IPv6 IPsec crypto maps are supported with IKE version v2 only.

Configuring Site-to-Site VPN

The following procedure describes how to configure a site-to-site VPN protecting IPv6 and IPv4 networks over an IPv6 IPsec tunnel.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand the **Site to Site** accordion.
3. Click **+** in the **IPSec Maps** table.
4. Enter the **IPV6 source network**, **IPV6 source prefix**, **IPV6 destination network**, **IPV6 destination prefix**, **Peer gateway v4 or v6** details in the **Create New IPSec** table.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure a site-to-site VPN protecting IPv6 networks over an IPv6 IPsec tunnel.

```
(host) [mynode] (config) #crypto-local ipsec-map <map name> <priority>
version v2
peer-ipv6 <IPv6 address>
vlan 1
src-net-ipv6 <IPv6 address> <Prefix length>
dst-net-ipv6 <IPv6 address> <Prefix length>
src-net <IPv4 address> <mask>
dst-net <IPv4 address> < mask>
```

Adding a New IPv6 Static Route

The following procedure describes how to to add a new IPv6 static route to an existing crypto-map.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces > IP Routes** tab.
2. Expand the **IP Routes** accordion.
3. Click **+** in the **IP Routes** table. Select **IPV6** for **IP version**.
4. Select **Using Forwarding Router Address** from the **Forwarding settings** drop-down list.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command adds a new IPv6 static route to an existing crypto-map.

```
(host) [md] (config) #ipv6 route <ipv6-network/prefix> ipsec <name>
```

Associating a PSK

The following procedure describes how to associate a PSK to the site-to-site crypto-map.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand the **Site to Site** accordion.
3. Click **+** in the **IPSec Maps** table and enter the following details in the **Create New Ipsec** table.
4. Select **Text-Based** or **Hex-based** from the **Representation type** drop-down list.
5. Enter the **IKE shared secret** and **Retype shared secret**.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands associate a pre-shared-key to the site-to-site crypto-map.

```
(host) [md] (config) #crypto-local isakmp key <key> addressv6 <IPv6 address> <prefix
length>
(host) [md] (config) #crypto-local isakmp key-hex <key> addressv6 <IPv6 address>
<prefix length>
```



Hex-based PSK is supported.

Associating a Certificate-Based Authentication

The following procedure describes how to associate a certificate based authentication to the site-to-site crypto-map:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand the **Site to Site** accordion.
3. Click **+** in the **IPSec Maps** table and enter the following details in the **Create New Ipsec** table.
   a. Select **Certificate** from the **Authentication method** drop-down list.
   b. Select values for **Server certificate** and **CA certificate** from the drop-down list.
   c. Enter a value for the **Peer certificate subject name** field.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands associate a certificate based authentication to the site-to-site crypto-map:

```
(host) [md] (config) #crypto-local ipsec-map <map name> <priority>
(host) [md] (config) #set ca-certificate <ca-certificate>
(host) [md] (config) #set server-certificate <server-certificate>
```



If you configure your Mobility Master to use IKEv2 and custom-installed certificates, you can optionally use Suite-B cryptographic algorithms for IPsec encryption. Select either **gcm-128 Use 128-bit AES-GCM Suite-B encryption** or **gcm-256 Use 256-bit AES-GCM Suite-B encryption**.

Monitoring and Managing Site-to-Site Settings

Execute the following command in the CLI to view the ISAKMP for an IPv6 peer:

```
(host) [md] #show crypto isakmp sa peer v6 2002::1
Initiator IP: 2002::1
Responder IP: 2002::3
Initiator: No
Initiator cookie:4ab9d9805eb16f73 Responder cookie:93f1c4dbec0ee92b
SA Creation Date: Fri Dec  4 23:14:33 2015
Life secs: 28800
Initiator Phase1 ID: 2002::1
Responder Phase1 ID: 2002::3
Exchange Type: IKE_SA (IKEV2)
Phase1 Transform:EncrAlg:AES128 HashAlg:HMAC_SHA1_96 DHGroup:2
Authentication Method: Pre Shared Key
IPSEC SA Rekey Number: 1
Ipsec-map name: default-local-master-ipsecmap2002::1
```

Execute the following command in the CLI to view the IPsec SA for an IPv6 peer:

```
(host) [md] #show crypto ipsec sa peer v6 2002::1
Initiator IP: 2002::1
Responder IP: 2002::3
Initiator: No
SA Creation Date: Sat Dec  5 00:50:01 2015
Life secs: 7200
Exchange Type: IKE_SA (IKEV2)
Phase2 Transform:Encryption Alg: 3DES  Authentication Alg: SHA1
Encapsulation Mode Tunnel
IP Compression Disabled
PFS: no
IN SPI: 1C514500, OUT SPI: 14F61800
Ipsec-map name: default-local-master-ipsecmap2002::1
Responder IP: 2002::3
```

Execute the following command in the CLI to view IKE transports:

```
(host) [md] #show crypto isakmp transports
transport 0x33cfb40 flags 0 refcnt 1
UDP-NATT Transport: fd 11   ikev2-id:0   src 1.1.1.10:4500 dst 1.1.1.4:4500
transport 0x2b3d660 flags 0 refcnt 1
UDP-500 Transport: fd 10   ikev2-id:0   src 1.1.1.10:500 dst 1.1.1.4:4500
transport 0x3292bb0 flags 0 refcnt 1
transport 0x298ea20 flags 1 refcnt 1
UDP-NATT Transport: fd 11   ikev2-id:0   src 0.0.0.0:4500 dst *:0
transport 0x298e940 flags 1 refcnt 1
UDP-500 Transport: fd 10   ikev2-id:0   src 0.0.0.0:500 dst
```

Execute the following command in the CLI to view IPv6 Switch address:

```
(host) [md] #show crypto isakmp stats
Switch IP                                          = 1.1.1.10
Main Mode Initiator exchanges started/completed    = 0/0
Main Mode Responder exchanges started/completed    = 0/0
Aggr Mode Initiator exchanges started/completed    = 0/0
Aggr Mode Responder exchanges started/completed    = 104034/0
Quick Mode Initiator exchanges started/completed   = 0/0
Quick Mode Responder exchanges started/completed   = 0/0
XAuth Type1 Responder exchanges started/completed  = 0/0
XAuth Type2 Responder exchanges started/completed  = 0/0
XAuth Authentication Pass/Fail                     = 0/0
Mode-Config Responder exchanges started/completed  = 0/0
Mode-Config Authentication Pass/Fail               = 0/0
XAuth Protocol Errors Bad-Packets/Quick-mode-fail   = 0/0
IP Pool        Alloc/Free/Free-NoSa    Alloc-Error/Free-Error = 0/0/0/0/0
IP External Pool  Alloc/Alloc-Error                = 0/0
Authentication State Errors  No-SA/No-Msg/No-Exch  = 0/0/0
Auth Msgs  Reqs/Rcvd/AP-Down/Idle-timeout/IP-down  = 0/0/0/0/0
Auth Msg Errors Not-Ready/Reqs-Throttled/IP-UP-err/Recv-err/Rcv-NoState  = 0/0/0/0/0
IKE->Auth Msgs  IP-up/IP-down                      = 0/0
Cert-Revocation Msgs  Reqs/Rcvd/Pass/Revoked       = 0/0/0/0
Cert-Revocation Msg Errors Reqs-Throttled/Send-err/Recv-err/Rcv-NoState = 0/0/0/0
UDB Msgs Reqs-Throttled/Req-sent/Req-send-errors/Resp-rcvd/Rcv-NoState = 0/0/0/0/0
ACR License Msgs Request/Delete/Req-errors/Resp-rcvd/Resp-error  = 0/0/0/0/0
Allow/Fail 0/0 Limit:1000
...
```

Execute the following commands in the CLI to clear IPsec and ISAKMP state security associations:

```
(host) [md] (config) #clear crypto isakmp sa peer v6 <>
(host) [md] (config) #clear crypto ipsec sa peer v6 <>
```

IP Compression Support for IPv6 Traffic Inside an IPsec Tunnel

Support for IP Compression is extended to IPv6 traffic inside an IPsec tunnel to minimize the size of the packets crossing a public network where ISP charges are calculated based on the number of bytes transferred.

IP Compression is supported for IPv6 traffic in an IPv4 IPsec Tunnel as well as IPv6 IPsec Tunnel.

Execute the following command in the CLI to enable or disable IP compression per crypto-map:

```
(host) [md] (config) #crypto-local ipsec-map test 9988
(host) [md] (config-submode)#ip-compression
(host) [md] (config-submode)#no ip-compression
```

Execute the following command in the CLI to verify if IP compression is enabled at the global level:

```
(host) [md] (config) #show crypto-local isakmp disable-ipcomp
IP Compression is Enabled
```

RADIUS Over IPv6

ArubaOS provides support for RADIUS authentication server over IPv6. You can configure an IPv6 host or specify an FQDN that can resolve to an IPv6 address for RADIUS authentication. The RADIUS server is in IPv4 mode by default. You must enable the RADIUS server in IPv6 mode to resolve the specified FQDN to IPv6 address.



You can only configure the global IPv6 address as the host for the Radius server in IPv6 mode.

The following procedure describes how to configure an IPv6 host for a RADIUS server:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication** > **Auth Servers** tab.
2. Select a server name from the **Server Groups** table.
3. Select a server name with the server type as **RADIUS** from **All Servers** table, to view the server details.
4. To enable the RADIUS server in IPv6 mode click **Enable IPv6** check box, in the **Server Options** table.
5. To configure an IPv6 host for the selected RADIUS server, specify an IPv6 address or an FQDN in the **IP address/hostname** field.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

To configure an IPv6 address for the NAS-IP:

1. Specify an IPv6 address in the **NAS IPv6** field.
2. Click **Submit**.
3. Click **Pending Changes**.
4. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands enable the `enable-ipv6` parameter:

```
(host) [mynode] (config) #aaa authentication-server radius IPv6
(host) [mynode] (RADIUS Server "IPv6") #enable-ipv6
```

Configure an IPv6 address as the host for RADIUS server using the following command:

```
(host) [mynode] (RADIUS Server "IPv6") #host <ipv6-address>
```

The <host> parameter can also be a FQDN that can resolve to an IPv6 address.



To resolve FQDN, you must configure the DNS server name using the `ip name-server <ip4addr>` command.

You can configure an IPv6 address for the NAS-IP parameter using the following CLI command:

```
(host) [mynode] (RADIUS Server "Ipv6") #nas-ip6 <IPv6 address>
```

You can configure an IPv6 address for the Source Interface parameter using the following CLI command:

```
(host) [mynode] (RADIUS Server "Ipv6") # source-interface vlan <vland-id> ip6addr
<ip6addr>
```

Use the following CLI command to configure an IPv6 address for the global NAS IP which the managed device uses to communicate with all the RADIUS servers:

```
(host) [mynode] (config) #ipv6 radius nas-ip6 <IPv6 address>
```

You can also configure an IPv6 global source-interface for all the RADIUS server requests using the following commands:

```
(host) [mynode] (config) #ipv6 radius source-interface loopback
(host) [mynode] (config) #ipv6 radius source-interface vlan <vlan-id> <ip6addr>
```

Radius Accounting for IPv6 Clients

Customers can now monitor bandwidth usage by clients or hosts with IPv6 addresses, over RADIUS protocol. The **Framed-IPv6-Address** attribute is used in accounting start, stop, and interim packets. A host can have multiple IPv6 addresses and all of them are tracked to check the usage, for billing purpose.

TACACS Over IPv6

ArubaOS provides support for TACACS authentication server over IPv6.

The following procedure describes how to configure the global IPv6 address.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** tab.
2. Select a server name from the **Server Groups** table.
3. Select a server name with the server type as **TACACS** from **All Servers** table, to view the server details.
4. To configure an IPv6 host for the selected server, specify an IPv6 address in the **IP address/hostname** field, in the **Server Options** table.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the global IPv6 address.

```
(host) [mynode] (config) #aaa authentication-server tacacs IPv6
(host) [mynode] (TACACS Server "IPv6") #host <ipv6-address>
```

DHCPv6 Server

The DHCPv6 server enables network administrators to configure stateful or stateless options and manage dynamic IPv6 users connecting to a network. You can also configure domain name server using DHCPv6.

You can configure IPv6 pools with various configurations such as lease duration, DNS server, vendor specific options, and user defined options using DHCPv6. You can also exclude IPv6 addresses from subnets. Managed Device IPv6 addresses, VLAN interface IPv6 addresses, and DNS server addresses are excluded from use by default.

Similar to DHCPv4, a DHCPv6 server pool is associated with a VLAN only through the IPv6 address configured in that VLAN interface. A VLAN interface can have a maximum of three global unicast addresses, but only one DHCPv6 pool.

DHCPv6 server supports stateless configuration of clients with options apart from the network addresses described in RFC 3736.

Points to Remember

- Similar to IPv4, the default router configuration is not required for IPv6 pools as IPv6-compliant routers will send RAs. The RA source address will be the default-gateway for the clients.
- ArubaOS does not support Hospitality feature on DHCPv6.

DHCP Lease Limit

From ArubaOS 8.2.0.0, the DHCP lease limits for 7000 Series controllers have changed.

The following table provides the maximum number of DHCP leases (both v4 and v6) supported per platform:



There is a new enforcement to the existing DHCP limit during configuration.

Table 27: *DHCP Lease Limits and Additional DHCP Scope*

Platform	Recommended DHCP Lease Limits in Previous Releases	DHCP Lease Limits in ArubaOS 8.2.0.0	Additional DHCP Scope Using CLI Command Option
Virtual Mobility Master	512	N/A	N/A
Virtual Mobility Controller - 32	512	N/A	N/A
Virtual Mobility Controller - 50	1024	N/A	N/A
Virtual Mobility Controller - 250	2048	N/A	N/A
Virtual Mobility Controller - 1K	4096	N/A	N/A
7005 Controller	512	1024	2048
7008 Controller	512	1024	2048
7010 Controller	1024	2048	4096
7024 Controller	1024	2048	2048
7030 Controller	2048	4096	4096
7205 Controller	4096	N/A	N/A
7210 Controller	5120	N/A	N/A
7220 Controller	10240	N/A	N/A
7240 Controller	15360	N/A	N/A
7240XM Controller	15360	N/A	N/A

For information on configuring or verifying additional DHCP scope, see the following topics:

- Configuring Additional DHCP Scope on page 156
- Verifying Additional DHCP Scope on page 156

Configuring Additional DHCP Scope

From ArubaOS 8.2.0.0, you can increase the DHCP lease limits to twice that of the user limits in 7005, 7008, and 7010 Controllers using the CLI. There is no WebUI option to configure the additional scope.

Execute the following command on a 7005, 7008, or 7010 controller to configure additional DHCP scope that is twice the user limit:

```
(host) (config) #ip dhcp increase-lease-limit
```

Verifying Additional DHCP Scope

Starting from ArubaOS 8.2.0.0, the output of the **show ip dhcp statistics** command is enhanced to show a warning if the DHCP lease limit of a 7005, 7008, or 7010 controller is increased beyond the user limit.

Execute the following **show** command on a 7005, 7008, or 7010 Controller to view the DHCP lease limit statistics:

```
(host) (config) #show ip dhcp statistics
DHCPv4 disabled; DHCPv6 disabled
DHCP Pools
----------
Network Name  Type  Active  Configured leases  Active leases  Free leases  Expired
leases
------------  ----  ------  -----------------  -------------  -----------  -----------
---
Abandoned leases
---------- ------
Current leases          0
Total leases            2048
```



DHCP lease limit increased beyond user limit. Some of the controller's services may be impacted.

To make a DHCPv6 pool active, ensure that the pool name is added in vlan interface.

Configuring DHCPv6 Server

You must enable the global DHCPv6 knob for the DHCPv6 functionality to be operational.

The following procedure describes how to configure DHCPv6 server.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > DHCP** tab.
2. Select the **IPv6 DHCP Server** check box to enable DHCPv6 globally.
3. If there are addresses that should not be assigned in the sub-network:
4. Under **IPv6 Excluded Address Range** table, click **+** to create a list of IPv6 excluded address.
5. Enter the excluded IPv6 address range in **IPv6 excluded range** and click **Submit**. The specified address range gets added to the **IPv6 Excluded Address** list box. The starting IP address in the **Exclude Address Range** should always contain a unique value, if the IP address is already present, then the existing IP address is replaced with a new one, and a warning is displayed.
6. Under **Pool Configuration**, click **+** to create a new DHCP server pool or click **Edit** to modify an existing DHCP server pool.



To enable the DHCPv6 Server functionality on an interface, select the **IP Interfaces** tab, edit the VLAN interface, and select a DHCP pool from the drop-down list under the **DHCP server** section. Ensure that the IP version of the VLAN interface is IPv6.

7. Select **IPv6** from the **IP version** drop-down list to create a DHCPv6 pool.
8. Enter a name in **Pool name** to configure an IPv6 pool name.
9. Enter an IPv6 address in **DNS servers** to configure an IPv6 DNS server.



To configure multiple DNS servers, enter the IPv6 addresses separated by space.

10. Enter a value in **Domain name** to configure the domain name.
11. Enter a value for **Preference**.
12. Enter the number of days, hours, minutes, and seconds in **Lease days, Lease hours, Lease minutes, and Lease seconds** to configure the lease time. The default value is 12 hours.
13. Specify an IPv6 prefix in **Network IP address** to configure an IPv6 network.
14. Enter the following details under **Option** to configure client specific DHCPv6 options.
    a. Under **Option**, click + to configure client specific DHCPv6 options.
    b. Specify the option code in **Option**.
    c. Select **IP** or **text** from the **IP/Text** drop-down list.
    d. Enter a value in **Value**. If you selected *IP* in *step b*, then you must enter a valid IPv6 address in this field.
    e. Click **OK**.
15. Click **Submit**.
16. Click **Pending Changes**.
17. In the **Pending Changes** window, select the check box and click **Deploy changes**.

To enable the DHCPv6 service you can use the following command:

```
(host) [md] (config)#service dhcpv6
To configure a domain name server, execute the following commands:
(host) [md] (config)#ipv6 dhcp pool <pool-name>
(host) [md] (config-dhcpv6)#dns-server <ipv6-address>
To configure a domain name, use the following command:
(host) [md] (config-dhcpv6)#domain-name <domain>
To configure DHCPv6 lease time, use the following command:
(host) [md] (config-dhcpv6)#lease <days> <hours> <minutes> <seconds>
The default value is 12 hours.
To configure a DHCP network, use the following command:
(host) [md] (config-dhcpv6)#network <network-prefix>
To configure a client specific option, use the following command:
(host) [md] (config-dhcpv6)#option <code> [ip <ipv6-address> | text <string>]
To configure DHCP server preference, use the following command:
(host) [md] (config-dhcpv6)#preference <value>
To enable DHCPv6 Server functionality on an interface, use the following command:
(host) [md] (config) #interface vlan <vlan-id>
(host) [md] (config-subif) #ipv6 dhcp server <pool-name>
```



The configured DHCPv6 pool subnet must match the interface prefix for DHCPv6 Server to be active.

To configure the IPv6 excluded address range for the DHCPv6 server, use the following command:

```
(host) [md] (config)#ipv6 dhcp excluded-address <low-address> [<high-address>]
```

You can view the DHCPv6 server settings, statistics, and binding information using the CLI.

To view the DHCPv6 database, use the following command:

```
(host) [md] #show ipv6 dhcp database
```

You can also view the DHCPv6 database for a specific pool, use the following command:

```
(host) [md] #show ipv6 dhcp database [pool <pool-name>]
(host) [md] #show ipv6 dhcp database pool DHCPv6
```

To view the DHCPv6 binding information, use the following command:

```
(host) [md]# show ipv6 dhcp binding
```

To clear all the DHCPv6 bindings, use the following command:

```
       (host) [md] # clear ipv6 dhcp binding
```

To view the DHCPv6 server statistics, use the following command:

```
(host) [md](config) #show ip dhcp statistics
```

To view the DHCPv6 active pools, use the following command:

```
(host) [md] #show ipv6 dhcp active-pools
```

Enabling DHCPv6 Relay

The DHCPv6 relay agent enhances the DHCP relay agent by providing support in an IPv6 network. The DHCPv6 relay agent passes messages between the DHCPv6 client and the DHCPv6 server, this is similar to the way DHCP relay agent supports an IPv4 network.

Starting from ArubaOS 8.2.0.0, you can configure DHCPv6 relay on a vlan interface.

The following procedure describes how to configure DHCPv6 relay.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Interfaces > VLANs** tab.
2. Select a VLAN from the **VLANs** table.
3. Under **Vlan Ids**, select the VLAN ID number.
4. Navigate to the **IPv6** tab for the selected VLAN ID.
5. Expand the **DHCP Server** accordion.
6. Select **DHCP Relay** from the **DHCP setting** drop-down list.
7. To add **DHCP helpers** for the VLAN, click **+** in the **DHCP helpers** accordion. Specify the following information in the pop-up window that appears:
8. The IPv6 address of the DHCPv6 server or any other relay agent to which the Mobility Master relays the DHCPv6 packets.
9. The source IPv6 address of the VLAN if more than one IPv6 address is configured.
10. Click **OK**.
11. Click **Submit**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Execute the following command to configure the DHCPv6 server relay agent.

```
(host) [mynode] (config) # ipv6 helper-address <address> source <srcaddr>
```

Execute the following command to view the helper address configured for a vlan.

```
(host) # show interface vlan <id>
```

Domain Name Server

Network devices on the internet use an IP address to route your request to the site you are trying to reach. Once you connect through a Domain Name Server it manages a database that maps domain names to IP addresses and routes your query to the next appropriate server.

Starting from ArubaOS 8.2.0.0, IPv6 DNS server configuration is also supported in addition to IPv4.

The following procedure describes how to configure Domain Name Server.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > General** tab.
2. Expand the **Domain Name System** accordion.
3. Select the **IPv4** and **IPv6** check boxes next to **Enable DNS name resolution** parameter to allow the user to enable or disable IPv4 and IPv6 DNS lookup.
4. To add a DNS Server, click **+** in the **DNS Servers** table. A pop-up window appears. Specify the appropriate values in **New DNS Server** window and click **OK**.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

To enable the IPv6 Domain Name Server hostname translation for clients, use the following command:

```
(host) [mynode] (config) #ipv6 domain lookup
```

To configure the IPv6 address of the domain name server, use the following command:

```
(host) [mynode] (config) #ipv6 name-server X:X:X:X::X
```

To view the domain name and server, use the following command:

```
(host) [mynode] #show ip domain-name
```

Redirect DNS Server

The redirect DNS server feature allows redirecting all DNS queries matching the corporate domain to a corporate DNS server. The DNS queries not matching the corporate domain are sent to the configured public DNS servers. You can configure corporate domain and corresponding DNS server in IPv4 or IPv6 profile. ArubaOS allows configuring up to three servers for a particular domain.

The following procedure describes how to configure Redirect DNS server.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > General** tab.
2. Expand the **Domain Name System** accordion.
3. Click the **Redirect DNS** toggle switch to enable this option.
4. To add a new Redirect DNS Server, click **+** in the **Domains to Redirect** table.
   a. In the **New Redirect DNS Server** window:
   b. Enter a **Domain name**.

c. Select the IP version.

d. Enter the **IPv4 or IPv6 address**.

5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

To enable Redirect DNS server, use the following command:

```
(host) [mynode] (config) #ip domain redirect
```

To configure the IPv4 redirect DNS server with a domain name, use the following command:

```
(host) [mynode] (config) #ip domain-redirect <domain-name> <server-ip>
```

For example:

```
(host) [mynode] (config) #ip domain-redirect arubanetworks.com 192.0.2.1
```

To configure the IPv6 redirect DNS server with domain name, use the following command:

```
(host) [mynode] (config) #ipv6 domain-redirect <domain-name> <server-ipv6>
```

For example:

```
(host) [mynode] (config) #ipv6 domain-redirect arubanetworks.com 2000::1
```

To view the domain name and server, use the following command:

```
(host) [mynode] #show ip domain-name
```

Understanding ArubaOS Supported Network Configuration for IPv6 Clients

ArubaOS provides wired or wireless clients using IPv6 addresses with services such as firewall functionality, layer-2 authentication, and, with the installation of the PEFNG, identity-based security. A managed device does not provide routing or NAT to IPv6 clients (see Understanding IPv6 Exceptions and Best Practices on page 166).

Supported Network Configuration

Clients can be wired or wireless and use IPv4 and/or IPv6 addresses. An external IPv6 router is recommended for a complete routing experience (dynamic routing). You can use the WebUI or CLI to display IPv6 client information.

A managed device can be configured with both IPv4 and IPv6 client addresses on the same VLAN.

Understanding the Network Connection Sequence for Windows IPv6 Clients

This section describes the network connection sequence for Windows Vista or XP clients that use IPv6 addresses, and the actions performed by the AP and the Managed Device.

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 managed device.
2. The managed device removes the 802.11 frame and creates an 802.3 frame for the Router Solicit message.

a. The managed device authenticates the user, applies firewall policies, and bridges the 802.3 frame to the IPv6 router.

b. The managed device creates entries in the user and session tables.

3. The IPv6 router responds with a Router Advertisement message.
4. The managed device applies firewall policies, then creates an 802.11 frame for the RA message. The managed device sends the RA through the GRE tunnel to the AP.
5. The IPv6 client sends a Neighbor Solicitation message.
6. The IPv6 router responds with a Neighbor Advertisement message.
7. If the DHCP is required to provide IPv6 addresses, the DHCPv6 process is started.
8. The IPv6 client sends data.

Understanding Authentication and Firewall Features that Support IPv6

This section describes ArubaOS features that support IPv6 clients.

Understanding Authentication

This release of ArubaOS only supports 802.1X authentication for IPv6 clients. You cannot configure Layer-3 authentications to authenticate IPv6 clients.

Table 28: *IPv6 Client Authentication*

Authentication Method	Supported for IPv6 Clients
802.1X	Yes
Stateful 802.1X (with non-Aruba APs)	Yes
Local database	Yes
Captive Portal	Yes
VPN	Yes
xSec	No (not tested)
MAC-based	Yes

You configure 802.1X authentication for IPv6 clients in the same way as for IPv4 client configurations. For more information about configuring 802.1X authentication on the Mobility Master, see 802.1X Authentication on page 268.



This release does not support authentication of management users on IPv6 clients.

Working with Firewall Features

If you installed a PEFNG license in the Mobility Master, you can configure firewall functions for IPv6 client traffic. While these firewall functions are identical to firewall functions for IPv4 clients, you need to

explicitly configure them for IPv6 traffic. For more information about firewall policies, see Understanding Global Firewall Parameters on page 509.



Voice-related and NAT firewall functions are not supported for IPv6 traffic.

Table 29: *IPv6 Firewall Parameters*

Parameter	Description
Monitor Ping Attack (per 30 seconds)	Number of ICMP pings per 30 second, which if exceeded, can indicate a DoS attack. Valid range is 1-16384 pings per 30 seconds. Recommended value is 120. Default: No default
Monitor TCP SYN Attack rate (per 30 seconds)	Number of TCP SYN messages per 30 second, which if exceeded, can indicate a DoS attack. Valid range is 1-16384 pings per 30 seconds. Recommended value is 960. Default: No default
Monitor IP Session Attack (per 30 seconds)	Number of TCP or UDP connection requests per 30 second, which if exceeded, can indicate a DoS attack. Valid range is 1-16384 requests per 30 seconds. Recommended value is 960. Default: No default
Deny Inter User Bridging	Prevents the forwarding of Layer-2 traffic between wired or wireless users. You can configure user role policies that prevent Layer-3 traffic between users or networks but this does not block Layer-2 traffic. This option can be used to prevent traffic, such as Appletalk or IPX, from being forwarded. Default: Disabled
Deny All IP Fragments	Drops all IP fragments. **NOTE:** Do not enable this option unless instructed to do so by an Aruba representative. Default: Disabled
Enforce TCP Handshake Before Allowing Data	Prevents data from passing between two clients until the three-way TCP handshake has been performed. This option should be disabled when you have mobile clients on the network, as enabling this option will cause mobility to fail. You can enable this option if there are no mobile clients on the network. Default: Disabled
Prohibit IP Spoofing	Enables detection of IP spoofing (where an intruder sends messages using the IP address of a trusted client). When you enable this option, IP and MAC addresses are checked for each ARP request or 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 Aruba representative. Default: Disabled
Session Mirror Destination	Destination (IPv4 address or managed device 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

Table 29: *IPv6 Firewall Parameters*

Parameter	Description
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 Aruba 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 Aruba representative, as doing so may create unnecessary overhead on the managed device. Default: Disabled (per-session logging is performed)
IPv6 Enable	Enables IPv6 globally.

The following procedure describes how to configure the firewall function:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > Firewall** tab.
2. Expand the **Global Setting** accordion.
3. Under the **IPv6** column, enter the following:
   a. Enter a value for **Monitor ping attack (per 30 sec).**
   b. Enter a value for **Monitor IP sessions attack (per 30 sec).**
   c. Enter a value for **Monitor TCP SYN attack rate (per 30 sec).**
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure firewall functions:

```
(host) [mynode] (config)#ipv6 firew all attack-rate ping 15
(host) [mynode] (config)#ipv6 firewall attack-rate session 25
(host) [mynode] (config)#ipv6 firewall session-idle-timeout 60
```

Understanding Firewall Policies

A user role, which determines a client's network privileges, is defined by one or more firewall policies. A firewall policy consists of rules that define the source, destination, and service type for specific traffic, and whether you want the managed device to permit or deny traffic that matches the rule.

You can configure firewall policies for IPv4 traffic or IPv6 traffic, and apply IPv4 and IPv6 firewall policies to the same user role. For example, if you have employees that use both IPv4 and IPv6 clients, you can configure both IPv4 and IPv6 firewall policies and apply them both to the "employee" user role.

The procedure to configure an IPv6 firewall policy rule is similar to configuring a firewall policy rule for IPv4 traffic, but with some differences. Table 30 describes the required and optional parameters for an IPv6 firewall policy rule.

Table 30: *IPv6 Firewall Policy Rule Parameters*

Parameter	Description
Source (required)	Source of the traffic: ■ **any**: Acts as a wildcard and applies to any source address. ■ **user**: This refers to traffic from the wireless client. ■ **host**: This refers to traffic from a specific host. When this option is chosen, you must configure the IPv6 address of the host. For example, 2002:d81f:f9f0:1000:c7e:5d61:585c:3ab. ■ **network**: This refers to a traffic that has a source IP from a subnet of IP addresses. When you chose this option, you must configure the IPv6 address and network mask of the subnet. For example, 2002:ac10:fe:: ffff:ffff:ffff::. ■ **alias**: This refers to using an alias for a host or network. **NOTE:** This release does not support IPv6 aliases. You cannot configure an alias for an IPv6 host or network.
Destination (required)	Destination of the traffic, which you can configure in the same manner as source.
Service (required)	**NOTE:** VoIP services are unavailable for IPv6 policies. Type of traffic: ■ **any**: This option specifies that this rule applies to any type of traffic. ■ **tcp**: Using this option, you configure a range of TCP ports to match the rule to be applied. ■ **udp**: Using this option, you configure a range of UDP ports to match the rule to be applied. ■ **service**: Using this option, you use one of the pre-defined services (common protocols such as HTTPS, HTTP, and others) as the protocol to match the rule to be applied. You can also specify a network service that you configure by navigating to the **Configuration > Advanced Services > Stateful Firewall > Network Services** page. ■ **protocol**: Using this option, you specify a different layer 4 protocol (other than TCP or UDP) by configuring the IP protocol value.
Action (required)	The action that you want the managed device to perform on a packet that matches the specified criteria. ■ **permit:** Permits traffic matching this rule. ■ **drop:** Drops packets matching this rule without any notification. **NOTE:** The only actions for IPv6 policy rules are permit or deny; in this release, the managed device cannot perform NAT or redirection on IPv6 packets. You can specify options such as logging, mirroring, or blacklisting (described below).
Log (optional)	Logs a match to this rule. This is recommended when a rule indicates a security breach, such as a data packet on a policy that is meant only to be used for voice calls.
Mirror (optional)	Mirrors session packets to a datapath or remote destination specified in the IPv6 firewall function (see Table 30). 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.

Table 30: *IPv6 Firewall Policy Rule Parameters*

Parameter	Description
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 ToS bits to be marked in the IP header of a packet matching this rule when it leaves the managed device.
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 managed device.

The following example creates a policy ipv6-web-only that allows only web (HTTP and HTTPS) access for IPv6 clients and assigns the policy to the user role “web-guest."



The user role web-guest can include both IPv6 and IPv4 policies, although this example only shows configuration of an IPv6 policy.

Creating an IPv6 Firewall Policy

The following procedure describes how to create an IPv6 firewall policy.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Policies** tab.
2. Click **+** to create a new policy.
3. Enter **ipv6-web-only** for the **Policy name**.
4. To configure a firewall policy, select **Session** for **Policy type**.
5. Click **Submit**.
6. Select the **ipv6-web-only** policy.
7. Click **+** in the **Policy > ipv6-web-only** rules table.
8. Select **Access Control** option in the **Rule Type** field and click **OK**.
9. Select **IPv6** from the **IP version** drop-down list.
10. Select **Network** from the **Source** drop-down list and enter the following values:
    a. For **IPv6 address**, enter **2002:d81f:f9f0:1000::**.
    b. For **IPv6 Netmask**, enter **64** as the prefix-length.
    c. For **Service/app**, select **Service** from the drop-down list.
    d. For **Service alias**, select **svc-http** from the drop-down list.
11. Click **Submit**.
12. Click **+ Policy > ipv6-web-only Rules table** to add a rule that allows HTTPS traffic.
13. Select **Access Control** option in the **Rule Type** field and click **OK**.
    a. Under **IP Version** column, select **IPv6.**
    b. Select **Network** from the **Source** drop-down list.
    c. For **IP**, enter **2002:d81f:f9f0:1000::**.
    d. For **Netmask**, enter **64** as the prefix-length.

e. Under **Service/app**, select **Service** from the drop-down list.

f. Select **svc-https** from the scrolling list.

14. Click **Submit**.



Rules can be reordered using the up and down arrow buttons provided for each rule.

15. Click **Pending Changes**.
16. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands create an IPv6 firewall policy.

```
(host) [md] (config)#ip access-list session ipv6-web-only
(host) [md] (config-submode)#ipv6 network 2002:d81f:f9f0:1000::/64 any svc-http permit
(host) [md] (config-submode)#ipv6 network 2002:d81f:f9f0:1000::/64 any svc-https
permit
```

Assigning an IPv6 Policy to a User Role

The following procedure describes how to assign an IPv6 policy.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Roles** tab.
2. Click **+** to create a new user role.
3. Enter **web-guest** in the **Name** field.
4. Click **Submit** .
5. Select **web-guest** role.
6. Click **Show Advanced View** .
7. Click **+** in **Roles > web-guest** table.
8. Select **Add an existing session policy** option, in the **New Policy** popup.
9. Select a policy from the **Policy name** drop-down list.
10. Click **Submit** .
11. Click **Pending Changes** .
12. In the **Pending Changes** window, select the check box and click **Deploy changes** .

The following CLI commands assign an IPv6 policy to a user role.

```
(host) [md] (config)#user-role web-guest
  (host) [md] (config-submode)#access-list session ipv6-web-only position 1
```

Understanding DHCPv6 Passthrough or Relay

The managed device forwards DHCPv6 requests from IPv6 clients to the external IPv6 router. On the external IPv6 router, you must configure the managed device's IP address as the DHCP relay. You do *not* need to configure an IP helper address on the managed device to forward DHCPv6 requests.

Understanding IPv6 Exceptions and Best Practices

The IPv6 best practices are provided below:

- Ensure that you enable IPv6 globally.
- The uplink port must be trusted. This is the same behavior as IPv4.

- Ensure that the `validuser` session ACL does not block IPv6 traffic.
- There must not be any ACLs that drop ICMPv6 or DHCPv6 traffic. It is acceptable to drop DHCPv6 traffic if the deployment uses SLAAC only.
- If an external device provides RA:
  - It is not recommended to advertise too many prefixes in RA.
  - The managed device supports a maximum of four IPv6 user entries in the user table. If a client uses more than four IPv6 addresses at a time, the user table is refreshed with the latest four active entries without disrupting the traffic flow. However, this may have some performance impact.
- Enable **BCMC Optimization** under interface VLAN to drop any random IPv6 multicast traffic. DHCPv6, ND, NS, and RA traffic are not dropped when you enable this option.



It is recommended to enable **BCMC Optimization** only if mDNS traffic is not used in the network, as mDNS traffic gets dropped if this option is enabled.

- While selecting a source address, the number of common bits between each source address in the list, is checked from the left most bit. This is followed by selection of the source address that has the maximum number of matching bits with the destination address. If more than one source addresses has the same number of matching bits with the destination address, the kernel selects that source address that is most recently configured on the system. It is essential that the administrator or user configures the network appropriately, if a particular VLAN interface needs to be selected as the source. For example, in case of 802.1X authentication the administrator or user can configure the source interface appropriately so that it is selected for authentication process. For more information on IPv6 source address selection, see **RFC 3848**.



Ensure that support for IPv6 Unique Local Address is added to enable configuring authentication-server hosts. .

ArubaOS does not support the following functions for IPv6 clients:

- The managed device offers limited routing services to IPv6 clients, so it is recommended to use an external IPv6 router for a complete routing experience (dynamic routing).
- VoIP ALG is not supported for IPv6 clients.
- IPv6 Auto configuration and IPv6 Neighbor Discovery mechanisms does not apply to IPv6 tunnels.
- Tunnel Encapsulation Limit, Tunnel-group, and MTU discovery options on IPv6 tunnels are not supported.
- When the **show upgrade managed-devices status copy all** command is executed after a managed device is upgraded, only the IPv4 address is displayed.
- IPv6 tunnel is not supported in tunnel-group. Hence, you cannot add Layer-2 or Layer-3 IPv6 GRE tunnels to a tunnel-group in both dual-stack and native IPv6 deployments.

OSPFv2 is a dynamic Interior IGP based on IETF RFC 2328. The OSPF uses the shortest or fastest routing path. Aruba's implementation of OSPFv2 allows Aruba Mobility Master and managed devices to deploy effectively in a Layer 3 topology. Aruba Mobility Master and managed devices can act as default gateway for all clients and forward user packets to the upstream router. The OSPF on the Mobility Master can be used to redistribute branch routes into corporate OSPF domain. The information on this chapter is in the following sections:

- Important Points to Remember on page 168
- Understanding OSPFv2 by Example using a WLAN Scenario on page 169
- Understanding OSPFv2 by Example using a Branch Scenario on page 170
- Configuring OSPF on page 171
- Exporting VPN Client Addresses to OSPF on page 171
- Sample Topology and Configuration on page 172

Important Points to Remember

OSPF is a robust routing protocol addressing various link types and deployment scenarios. The Aruba implementation applies to two main use cases; WLAN Scenarios and Branch Scenario.

- OSPF is disabled by default.
- Aruba Mobility Master supports only one OSPF instance.
- Convergence takes between 5 and 15 seconds.
- All area types are supported.
- Multiple configured areas are supported.
- An Aruba Mobility Master can act as an ABR.
- OSPF supports VLAN and GRE tunnel interfaces.
- To run OSPF over IPsec tunnels, a Layer 3 GRE tunnel is configured between two routers with GRE destination addresses as the inner address of the IPsec tunnel. OSPF is enabled on the Layer 3 GRE tunnel interface, and all of the OSPF control packets undergo GRE encapsulation before entering the IPsec tunnels.
- The default MTU value for a Layer 3 GRE tunnel in an Aruba Mobility Master is 1100. When running OSPF over a GRE tunnel between an Aruba Mobility Master and another vendor's router, the MTU values must be the same on both sides of the GRE tunnel.

The following table provides information on the maximum OSPF routes supported for various platforms:

Table 31: *Maximum OSPF Routes*

Platform	Branches	Routes
7005	4K	4K

Platform	Branches	Routes
7008	4K	4K
7210	8K	8K
7220	16K	16K
7240	32K	32K

Below are some guidelines regarding deployment and topology for this release of OSPFv2:

- In the WLAN scenario upstream router, configure only the interface connected to the stand-alone controller or the managed device in the same area. This will minimize the number of local subnet addresses advertised by the upstream router to the stand-alone controller or the managed device.
- Use the upstream router as the designated router for the link or interface between the stand-alone controller or the managed device and the upstream router.
- The default MTU value for a Layer 3 GRE tunnel in an Aruba Mobility Master, managed device or stand-alone controller is 1100. When running OSPF over a GRE tunnel between an Aruba device and another vendor's router, the MTU values must be the same on both sides of the GRE tunnel.
- Do not enable OSPF on any uplink or WAN interfaces on the managed device. Enable OSPF only on the Layer 3 GRE tunnel connecting the Mobility Master.
- Use only one physical port in the uplink VLAN interface that is connecting to the upstream router. This will prevent broadcasting the protocol PDUs to other ports and hence limit the number of adjacencies on the uplink interface to only one.

Understanding OSPFv2 by Example using a WLAN Scenario

In the WLAN scenario, the Aruba Mobility Master acts as a default gateway for all the clients, and talks to one or two upstream routers for redundancy. Mobility Master advertises all the user subnet addresses as stub addresses to the routers via LSAs.



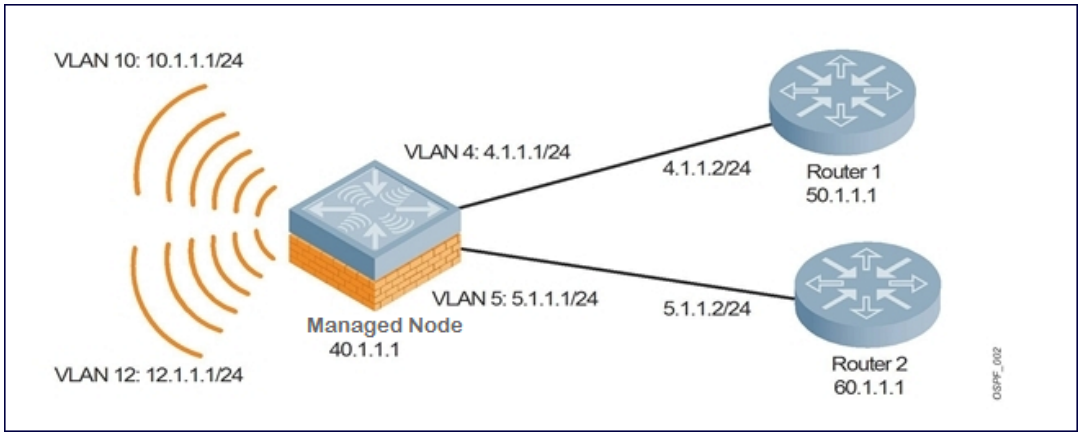
Totally stub areas see only default route and to the areas themselves.

WLAN Topology

Mobility Master (Figure 17) is configured with VLAN 10 and VLAN 12 as user VLANs. These VLANs have clients on the subnets, and Mobility Master 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:

- Aruba managed device— 40.1.1.1
- Router 1— 50.1.1.1
- Router 2— 60.1.1.1

Figure 17 *WLAN OSPF Topology*



WLAN Routing Table

View the Mobility Master routing table using the **show ip route** command:

```
(host) [mynode] #show ip route

Codes: C - connected, O - OSPF, R - RIP, S - static
       M - mgmt, U - route usable, * - candidate default
```

Below is the routing table for Router 1:

```
(router1) #show ip route

O    10.1.1.0/24  [1/0] via 4.1.1.1
O    12.1.1.0/24  [1/0] via 4.1.1.1
```

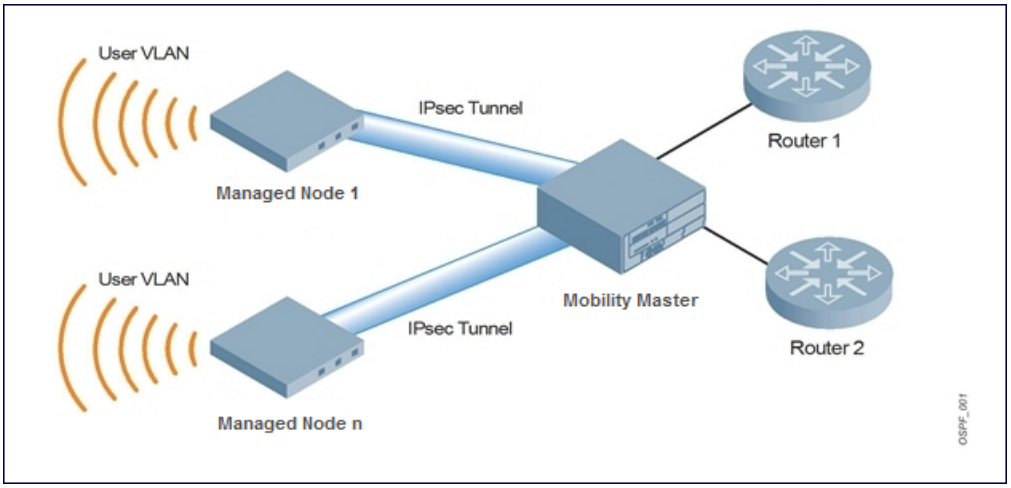
Understanding OSPFv2 by Example using a Branch Scenario

The branch office scenario has a number of remote branch offices with managed devices talking to a central office via a Mobility Master using site-to-site VPN tunnels or IPsec tunnels. The central office Mobility Master is in turn talking to the upstream routers (see Figure 18). 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 managed device in the branch office to reach the corporate subnets.

Branch Topology

All the OSPF control packets exchanged between the managed devices and Mobility Master undergo GRE encapsulation before entering the IPsec tunnels. The managed devices in the branch offices advertise all the user subnet addresses to Mobility Master as stub addresses in router LSA. Mobility Master in turn forwards those router LSAs to the upstream routers.

Figure 18 *Branch OSPF Topology*



Branch Routing Table

View the branch office managed device routing table using the **show ip route** command:

```
(host) [md] #show ip route

Codes: C - connected, O - OSPF, R - RIP, S - static
       M - mgmt, U - route usable, * - candidate default
```

The routing table for Mobility Master is below:

```
(host) [mynode] #show ip route

Gateway of last resort is 4.1.1.2 to network 0.0.0.0

O*    0.0.0.0/0  [1/0] via 4.1.1.2*
O    14.1.1.0/24  [1/0] via 30.1.1.1*
O    15.1.1.0/24  [1/0] via 30.1.1.1*
C    4.1.1.0 is directly connected, VLAN4
C    5.1.1.0 is directly connected, VLAN5
```

Configuring OSPF

To configure general OSPF settings from the OSPF tab, perform the following steps:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration** > **Interfaces** > **OSPF** tab.
2. To enable OSPF, Select the **Enable OSPF** toggle switch.
3. Configure the other OSPF interface settings in the respective fields.
4. To add an OSPF area, click the **+** icon in the **Area** table and specify the appropriate values.
5. To add an excluded subnet in the **Excluded Subnet** table, click the **+** icon and specify the appropriate values.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Exporting VPN Client Addresses to OSPF

You can configure VPN client addresses so that they can be exported to OSPF and be advertised as host routes (/32). Exporting applies to any VPN client address regardless of how it is assigned.

In the WebUI

To export a VPN client address to OSPF using the WebUI:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration** > **System** > **Profiles** > **Wireless LAN** > **VPN Authentication** > **default** page.

   For an Instant AP, Remote AP and Campus AP, you can edit the respective default profiles (default-iap, default-rap, and default-cap)

2. (Optional) Select the **Export VPN IP address as a route** check box. Regardless of how an authentication server is contacted, selecting this option causes any VPN client address to be exported to OSPF using IPC. Note that the Framed-IP-Address attribute is assigned the IP address as long as any server returns the attribute. The Framed-IP-Address value always has a higher priority than the local address pool.
3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

In the CLI

To export a VPN client address to OSPF using CLI:

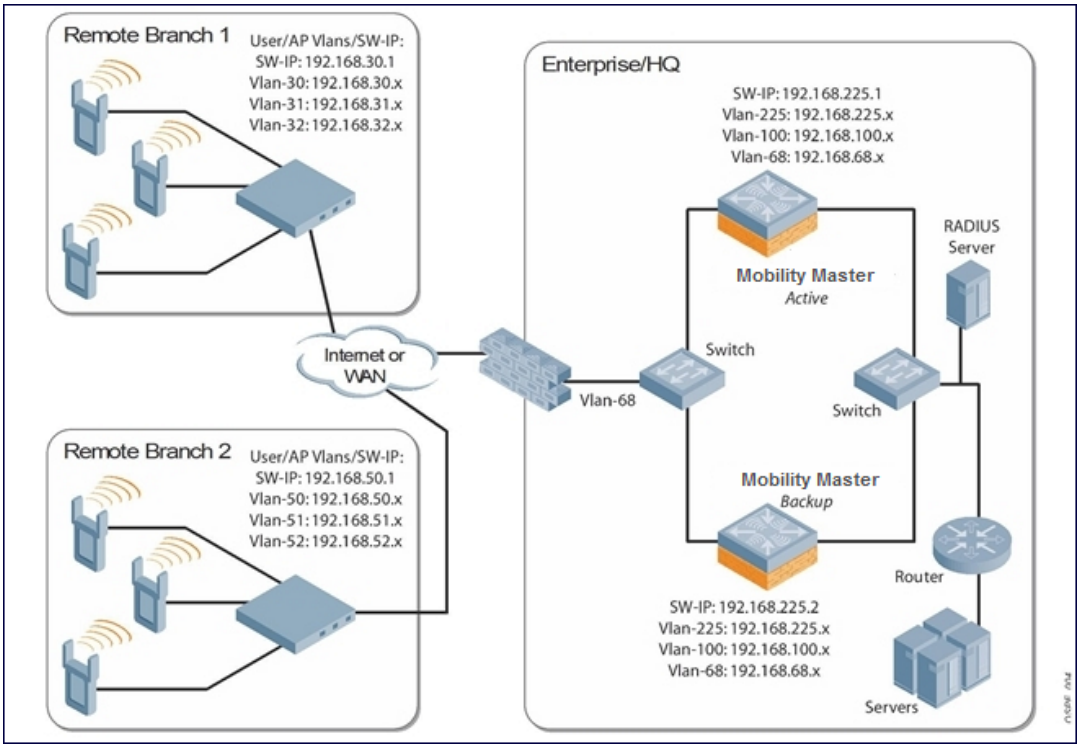
```
(host) [mynode] (config) #aaa authentication vpn default
(host) [mynode] (VPN Authentication Profile "default") #
(host) [mynode] (VPN Authentication Profile "default") # export-route
```

Use the **show ip ospf** database command to show LSA types that are generated.

Sample Topology and Configuration

The figure below displays a sample OSPF topology followed by sample configurations of the Remote Branch 1, Remote Branch 2, and the Central Office Mobility Master (Active and Backup).

Figure 19 *Sample OSPF Topology*



Remote Branch 1

```
controller-ip vlan 30
vlan 16
vlan 30
vlan 31
vlan 32
interface gigabitethernet  0/0/1
        description "GE0/0/1"
        trusted
        switchport access vlan 16
!
interface gigabitethernet  0/0/2
        description "GE0/0/2"
        trusted
        switchport access vlan 30
!

interface gigabitethernet  0/0/3
        description "GE0/0/3"
        trusted
        switchport access vlan 31
!
interface gigabitethernet  0/0/4
        description "GE0/0/4"
        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  0/0/1
        description "GE0/0/1"
        trusted
        switchport access vlan 20
!
interface gigabitethernet  0/0/2
        description "GE0/0/2"
        trusted
        switchport access vlan 50
!
interface gigabitethernet  0/0/3
        description "GE0/0/3"
        trusted
        switchport access vlan 51
!
interface gigabitethernet  0/0/4
        description "GE10/0/4"
        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
```

Mobility Master—Active

```
localip 0.0.0.0 ipsec db947e8d1b383813a4070ab0799fa6246b80fc5cfcc3268f
controller-ip vlan 225
vlan 68
vlan 100
vlan 225
!
interface gigabitethernet  0/0/1
        description "GE0/0/1"
        trusted
        switchport access vlan 225
!
interface gigabitethernet  0/0/2
        description "GE0/0/2"
        trusted
        switchport access vlan 100
!
interface gigabitethernet  10/0/31
        description "GE0/0/4"
        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
!
```

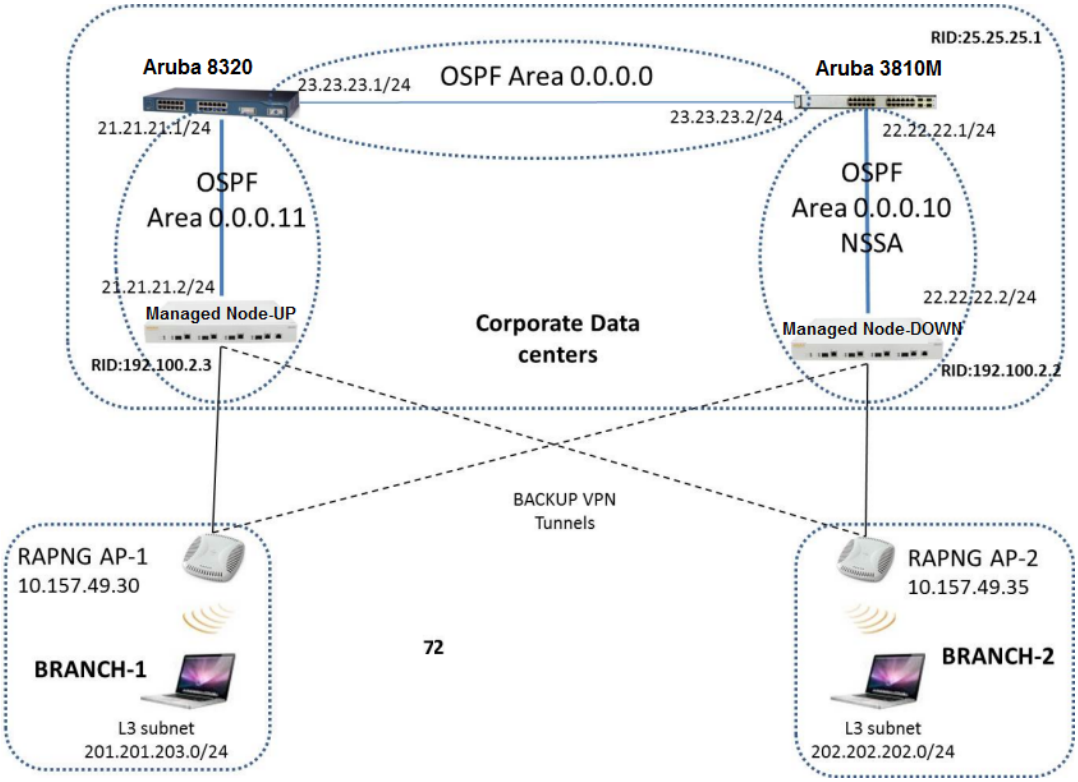
Mobility Master—Backup

```
localip 0.0.0.0 ipsec db947e8d1b383813a4070ab0799fa6246b80fc5cfcc3268f
controller-ip vlan 225
!
interface gigabitethernet  0/0/1
        description "GE0/0/1"
        trusted
        switchport access vlan 225
!
interface gigabitethernet  0/0/2
        description "GE0/0/2"
        trusted
        switchport access vlan 100
!
interface gigabitethernet  0/0/31
        description "GE0/0/3"
        trusted
        switchport access vlan 68
!
interface vlan 68
        ip address 192.168.68.221 255.255.255.224
!
interface vlan 100
        ip address 192.168.100.5 255.255.255.0
```

```
!
interface vlan 225
        ip address 192.168.225.1 255.255.255.0
!
interface tunnel 2003
        description "Tunnel Interface"
        ip address 2.1.0.3 255.0.0.0
        tunnel source 192.168.225.1
        tunnel destination 192.168.30.1
        trusted
        ip ospf area 10.10.10.10
!
interface tunnel 2005
        description "Tunnel Interface"
        ip address 2.1.0.5 255.0.0.0
        tunnel source 192.168.225.1
        tunnel destination 192.168.50.1
        trusted
        ip ospf area 10.10.10.10
!
master-redundancy
  master-vrrp 2
  peer-ip-address 192.168.68.220 ipsec password123
!
vrrp 1
  priority 99
  authentication password123
  ip address 192.168.68.217
  vlan 68
  tracking vlan 68 sub 40
  tracking vlan 100 sub 40
  tracking vlan 225 sub 40
  no shutdown
!
vrrp 2
  priority 99
  ip address 192.168.225.9
  vlan 225
  tracking vlan 68 sub 40
  tracking vlan 100 sub 40
  tracking vlan 225 sub 40
  no shutdown
!
ip default-gateway 192.168.68.1
ip route 192.168.0.0 255.255.0.0 null 0
!
router ospf
router ospf router-id 192.168.225.1
router ospf area 10.10.10.10 stub
router ospf redistribute vlan 100,225
!
```

The following figure displays how the managed device is configured for Instant AP VPN for different OSPF cases.

Figure 20 *Managed Device Not-So-Stubby-Area and Normal Area*



Topology

- Area-10 is Not-So-Stubby Area
- Area-11 is Normal area.
- RAPNG AP-1 is configured to have a UP managed device as its primary managed device and a DOWN as secondary managed device.
- RAPNG AP-2 is configured to have a DOWN as its primary managed device and a UP as secondary managed device.
- RAPNG AP-1 is configured to have a 201.201.203.0/24 L3-distributed network.
- RAPNG AP-2 is configured to have a 202.202.202.0/24 L3-distributed network.

Observation

- UP managed device will send Type-5 LSA (External LSA) of VPN route 201.201.203.0/24 to it's upstream router, Aruba 3810M.
- DOWN managed device will send Type-7 LSA (NSSA) of VPN route 202.202.202.0/24 to it's upstream router, Aruba 8320.
- UP managed device will send a Type-4 asbr-summary LSA.

Configuring UP Managed Device

```
interface vlan 21
ip address 21.21.21.2 255.255.255.0
ip ospf area 0.0.0.11
!
router ospf
router ospf area 0.0.0.11
```

```
router ospf redistribute rapng-vpn
!
```

The following commands display the configuration and run time protocol details on the UP managed device:

```
(host) [mynode]#show ip route
Codes: C - connected, O - OSPF, R - RIP, S - static
M - mgmt, U - route usable, * - candidate default, V - RAPNG VPN
Gateway of last resort is Imported from DHCP to network 0.0.0.0 at cost 10
Gateway of last resort is Imported from CELL to network 0.0.0.0 at cost 10
Gateway of last resort is Imported from PPPOE to network 0.0.0.0 at cost 10
Gateway of last resort is 10.15.231.185 to network 0.0.0.0 at cost 1
S*    0.0.0.0/0  [1/0] via 10.15.231.185*
O    10.15.228.0/27 [333/0] via 21.21.21.1*
O    12.12.12.0/25 [0/0] via 21.21.21.1*
O    22.22.22.0/24 [3/0] via 21.21.21.1*
O    23.23.23.0/24 [2/0] via 21.21.21.1*
O    25.25.25.0/24 [333/0] via 21.21.21.1*
S    192.100.3.0/24 [1/0] via 192.100.2.1*
S    192.100.4.0/24 [1/0] via 192.100.2.1*
S    192.100.5.0/24 [1/0] via 192.100.2.1*
S    192.100.6.0/24 [1/0] via 192.100.2.1*
S    192.100.7.0/24 [1/0] via 192.100.2.1*
S    192.100.8.0/24 [1/0] via 192.100.2.1*
S    192.100.9.0/24 [1/0] via 192.100.2.1*
S    192.100.10.0/24 [1/0] via 192.100.2.1*
S    192.100.11.0/24 [1/0] via 192.100.2.1*
S    192.100.12.0/24 [1/0] via 192.100.2.1*
S    192.100.13.0/24 [1/0] via 192.100.2.1*
S    192.100.14.0/24 [1/0] via 192.100.2.1*
S    192.168.1.0/24 [1/0] via 192.100.2.1*
S    192.169.1.0/24 [1/0] via 192.100.2.1*
S    192.170.1.0/24 [1/0] via 192.100.2.1*
S    192.171.1.0/24 [1/0] via 192.100.2.1*
S    192.172.1.0/24 [1/0] via 192.100.2.1*
S    192.173.1.0/24 [1/0] via 192.100.2.1*
S    192.174.1.0/24 [1/0] via 192.100.2.1*
S    192.175.1.0/24 [1/0] via 192.100.2.1*
S    192.176.1.0/24 [1/0] via 192.100.2.1*
S    192.177.1.0/24 [1/0] via 192.100.2.1*
S    192.178.1.0/24 [1/0] via 192.100.2.1*
S    192.179.1.0/24 [1/0] via 192.100.2.1*
V    201.201.203.0/26 [10/0] ipsec map
O    202.202.202.0/29 [0/0] via 21.21.21.1*
C    192.100.2.0/24 is directly connected, VLAN2
C    10.15.231.184/29 is directly connected, VLAN1
C    172.16.0.0/24 is directly connected, VLAN3
C    21.21.21.0/24 is directly connected, VLAN21
C    5.5.0.2/32 is an ipsec map 10.15.149.30-5.5.0.2

(host) [mynode]#show ip ospf database
OSPF Database Table
-------------------
Area ID    LSA Type        Link ID         Adv Router      Age    Seq#
Checksum
-------    --------        -------         ----------      ---    ----         -------
-
0.0.0.11   ROUTER          21.21.21.1      21.21.21.1      178    0x80000017   0xca50
0.0.0.11   ROUTER          192.100.2.3     192.100.2.3     1406   0x80000007   0x2253
0.0.0.11   NETWORK         21.21.21.1      21.21.21.1      178    0x80000003   0xdf6d
0.0.0.11   IPNET_SUMMARY   22.22.22.0      21.21.21.1      178    0x80000003   0x7e38
0.0.0.11   IPNET_SUMMARY   23.23.23.0      21.21.21.1      178    0x80000003   0x5064
0.0.0.11   ASBR_SUMMARY    25.25.25.1      21.21.21.1      178    0x80000003   0xefbc
```

```
0.0.0.11    ASBR_SUMMARY    192.100.2.3     192.100.2.3     1412    0x80000002   0xa85d
N/A         AS_EXTERNAL     10.15.228.0     25.25.25.1      1014    0x8000000e   0xea43
N/A         AS_EXTERNAL     12.12.12.0      25.25.25.1      268     0x80000003   0x433a
N/A         AS_EXTERNAL     25.25.25.0      25.25.25.1      1761    0x80000005   0x3d8d
N/A         AS_EXTERNAL     201.201.203.0   10.15.231.186   3600    0x80000001   0x6690
N/A         AS_EXTERNAL     201.201.203.0   192.100.2.3     1104    0x80000002   0xe4a2
N/A         AS_EXTERNAL     202.202.202.0   25.25.25.1      268     0x80000003   0x4385

(host) [mynode]#show ip ospf neighbor

OSPF Neighbor Table
-------------------
Neighbor ID  Pri  State     Address      Interface
-----------  ---  -----     -------      ---------
21.21.21.1   1    FULL/DR   21.21.21.1   Vlan
```

Configuring DOWN Managed Device

```
interface vlan 22
ip address 22.22.22.2 255.255.255.0
ip ospf area 0.0.0.10
!
router ospf
router ospf area 0.0.0.10 nssa
router ospf redistribute rapng-vpn
!
```

The following commands display the configuration and run time protocol details on the DOWN managed device:

```
(host) [mynode]#show ip route

Codes: C - connected, O - OSPF, R - RIP, S - static
M - mgmt, U - route usable, * - candidate default, V - RAPNG VPN
Gateway of last resort is Imported from DHCP to network 0.0.0.0 at cost 10
Gateway of last resort is Imported from CELL to network 0.0.0.0 at cost 10
Gateway of last resort is Imported from PPPOE to network 0.0.0.0 at cost 10
O    0.0.0.0/0  [1/0] via 22.22.22.1*
S    10.0.0.0/8  [1/0] via 10.15.231.177*
O    10.15.228.0/27 [333/0] via 22.22.22.1*
V    12.12.12.0/25 [10/0] ipsec map
O    21.21.21.0/24 [3/0] via 22.22.22.1*
O    23.23.23.0/24 [2/0] via 22.22.22.1*
O    25.25.25.0/24 [333/0] via 22.22.22.1*
V    202.202.202.0/29 [10/0] ipsec map
C    192.100.2.0/24 is directly connected, VLAN2
C    10.15.231.176/29 is directly connected, VLAN1
C    22.22.22.0/24 is directly connected, VLAN22
C    4.4.0.2/32 is an ipsec map 10.15.149.35-4.4.0.2
C    4.4.0.1/32 is an ipsec map 10.17.87.126-4.4.0.1

(host) [mynode]#show ip ospf neighbor
OSPF Neighbor Table
-------------------
Neighbor ID  Pri  State      Address      Interface
-----------  ---  -----      -------      ---------
25.25.25.1   1    FULL/BDR   22.22.22.1   Vlan 22

(host) [mynode]#show ip ospf database
OSPF Database Table
-------------------
Area ID    LSA Type        Link ID         Adv Router    Age    Seq#         Checksum
-------    --------        -------         ----------    ---    ----         --------
```

```
0.0.0.10    ROUTER          25.25.25.1      25.25.25.1    1736   0x80000021   0xb732
0.0.0.10    ROUTER          192.100.2.2     192.100.2.2   500    0x80000005   0x9ad9
0.0.0.10    NETWORK         22.22.22.2      192.100.2.2   500    0x80000004   0x8aeb
0.0.0.10    IPNET_SUMMARY   21.21.21.0      25.25.25.1    1990   0x80000003   0xe7bf
0.0.0.10    IPNET_SUMMARY   23.23.23.0      25.25.25.1    1990   0x80000003   0x950d
0.0.0.10    NSSA            0.0.0.0         25.25.25.1    725    0x80000002   0xaab9
0.0.0.10    NSSA            10.15.228.0     25.25.25.1    1228   0x80000010   0xca5f
0.0.0.10    NSSA            12.12.12.0      192.100.2.2   352    0x80000005   0xe8cb
0.0.0.10    NSSA            25.25.25.0      25.25.25.1    1485   0x80000006   0x1fa8
0.0.0.10    NSSA            202.202.202.0   192.100.2.2   352    0x80000005   0xe817
N/A         AS_EXTERNAL     12.12.12.0      192.100.2.2   352    0x80000005   0x28d8
N/A         AS_EXTERNAL     202.202.202.0   192.100.2.2   352    0x80000005   0x2824
```

Viewing the Status of Instant AP VPN

You can view the status of an Instant AP VPN for RAPNG AP-1 and RAPNG AP-3, using the following commands:

RAPNG AP-1

```
(host) [mynode]# show  vpn status

profile name:default
-------------------------------------------------
current using tunnel                            :primary tunnel
ipsec is preempt status                         :disable
ipsec is fast failover status                   :disable
ipsec hold on period                            :600
ipsec tunnel monitor frequency (seconds/packet) :5
ipsec tunnel monitor timeout by lost packet cnt :2
ipsec     primary tunnel crypto type            :Cert
ipsec     primary tunnel peer address           :10.15.231.186
ipsec     primary tunnel peer tunnel ip         :192.100.2.3
ipsec     primary tunnel ap tunnel ip           :5.5.0.2
ipsec     primary tunnel current sm status      :Up
ipsec     primary tunnel tunnel status          :Up
ipsec     primary tunnel tunnel retry times     :2
ipsec     primary tunnel tunnel uptime          :1 hour 24 minutes 50 seconds
ipsec      backup tunnel crypto type            :Cert
ipsec      backup tunnel peer address           :10.15.231.178
ipsec      backup tunnel peer tunnel ip         :0.0.0.0
ipsec      backup tunnel ap tunnel ip           :0.0.0.0
ipsec      backup tunnel current sm status      :Init
ipsec      backup tunnel tunnel status          :Down
ipsec      backup tunnel tunnel retry times     :0
ipsec      backup tunnel tunnel uptime          :0
```

(host)# show datapath route

```
Route Table Entries
-------------------
Flags: L - Local, P - Permanent,  T - Tunnel, I - IPsec, M - Mobile, A - ARP, D - Drop
IP              Mask            Gateway        Cost  VLAN  Flags
--------------  --------------  --------------  ----  ----  -----
0.0.0.0         0.0.0.0         10.15.149.25       0     0
0.0.0.0         128.0.0.0       192.100.2.3        0     0  T
128.0.0.0       128.0.0.0       192.100.2.3        0     0  T
192.168.10.0    255.255.254.0   192.168.10.1       0  3333  D
201.201.203.0   255.255.255.192 0.0.0.0            0   103  LP
10.15.149.24    255.255.255.248 10.15.149.30       0     1  L
10.15.231.186   255.255.255.255 10.15.149.25       0     0
Route Cache Entries
-------------------
Flags: L - local, P - Permanent,  T - Tunnel, I - IPsec, M - Mobile, A - ARP, D - Drop
```

```
IP              MAC              VLAN      Flags
--------------  -----------------  -----------  -----
202.202.202.6    00:00:00:00:00:00            0  T
192.100.2.3      00:00:00:00:00:00            0  PT
192.168.10.51    10:40:F3:98:80:94            1  PA
192.168.10.1     00:24:6C:C9:27:A3         3333  LP
201.201.203.8    00:26:C6:52:6B:14          103
201.201.203.1    00:24:6C:C9:27:A3          103  LP
10.1.1.50        00:00:00:00:00:00            0  T
```

RAPNG AP-3

```
(host) [mynode]# show vpn status
profile name:default
-------------------------------------------------
current using tunnel                            :primary tunnel
ipsec is preempt status                         :disable
ipsec is fast failover status                   :disable
ipsec hold on period                            :600
ipsec tunnel monitor frequency (seconds/packet) :5
ipsec tunnel monitor timeout by lost packet cnt :2
ipsec     primary tunnel crypto type            :Cert
ipsec     primary tunnel peer address           :10.15.231.178
ipsec     primary tunnel peer tunnel ip         :192.100.2.2
ipsec     primary tunnel ap tunnel ip           :4.4.0.2
ipsec     primary tunnel current sm status      :Up
ipsec     primary tunnel tunnel status          :Up
ipsec     primary tunnel tunnel retry times     :13
ipsec     primary tunnel tunnel uptime          :1 hour 55 minutes 6 seconds
ipsec      backup tunnel crypto type            :Cert
ipsec      backup tunnel peer address           :10.15.231.186
ipsec      backup tunnel peer tunnel ip         :0.0.0.0
ipsec      backup tunnel ap tunnel ip           :0.0.0.0
ipsec      backup tunnel current sm status      :Init
ipsec      backup tunnel tunnel status          :Down
ipsec      backup tunnel tunnel retry times     :0
ipsec      backup tunnel tunnel uptime          :0

(host) [mynode]# show datapath route

Route Table Entries
-------------------
Flags: L - Local, P - Permanent,  T - Tunnel, I - IPsec, M - Mobile, A - ARP, D - Drop
IP              Mask            Gateway        Cost  VLAN  Flags
--------------  --------------  ---------------  ----  ----  -----
0.0.0.0          0.0.0.0         10.15.149.33        0     0
0.0.0.0          128.0.0.0       192.100.2.2         0     0  T
128.0.0.0        128.0.0.0       192.100.2.2         0     0  T
192.168.10.0     255.255.254.0   192.168.10.1        0  3333  D
10.15.149.32     255.255.255.248 10.15.149.35        0     1  L
202.202.202.0    255.255.255.248 0.0.0.0             0   203  LP
10.15.231.178    255.255.255.255 10.15.149.33        0     0
Route Cache Entries
-------------------
Flags: L - local, P - Permanent,  T - Tunnel, I - IPsec, M - Mobile, A - ARP, D - Drop
IP              MAC              VLAN      Flags
--------------  -----------------  -----------  -----
202.202.202.1    00:24:6C:C0:41:F2          203  LP
202.202.202.6    08:ED:B9:E1:51:7B          203
192.100.2.2      00:00:00:00:00:00            0  PT
192.168.10.1     00:24:6C:C0:41:F2         3333  LP
201.201.203.8    00:00:00:00:00:00            0  T
10.1.1.50        00:00:00:00:00:00            0  T
192.168.11.7     00:26:C6:52:6B:14            1  PA
```

```
4.4.0.2          00:24:6C:C0:41:F2           1  LP
10.13.6.110      00:00:00:00:00:00           0  T
10.15.149.38     00:24:6C:C9:27:CC           1  A
10.15.149.35     00:24:6C:C0:41:F2           1  LP
10.15.149.33     00:0B:86:40:93:00           1  A
```

This chapter describes how to configure a tunneled node, also known as a wired tunneled node. Tunneled nodes provide access and security using an overlay architecture.

This chapter describes the following topics:

- Understanding Tunneled Node Configuration on page 184
- Configuration Procedures on page 185

Understanding Tunneled Node Configuration

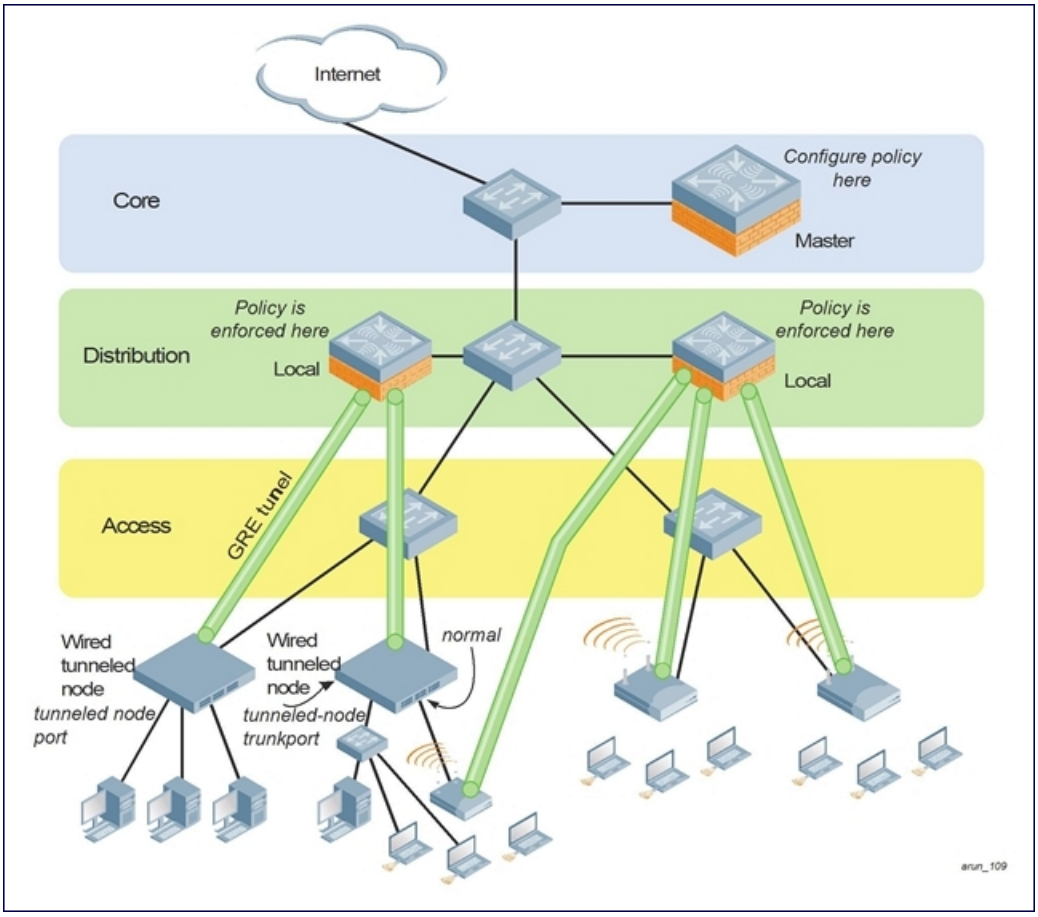
The tunneled node connects to one or more client devices at the edge of the network and then establishes a secure GRE tunnel to the controlling concentrator server. This approach allows the managed device to support all the centralized security features, like 802.1X authentication, captive-portal authentication, and stateful firewall. A tunneled node is required to handle only the physical connection to clients.

To support the wired concentrator, the managed device must have a license to terminate APs, no other configuration is required. To configure the tunneled node, specify the IP address of the managed device and identify the ports that should be used as active tunneled node ports. Tunnels are established between the managed device and each active tunneled node port on the tunneled node. All tunneled node units must run the same version of ArubaOS. The tunneled node port can also be configured as a trunk port. This allows customers to have multiple clients on different VLANs that come through the trunk port instead of having clients on a single VLAN.

Figure 21 shows how the tunneled node fits into network operations. Traffic moves through GRE tunnels between the active tunneled node ports and the managed device. Policies are configured on the managed device and can be enforced on the same managed device or on different systems.

On the managed device, you can assign the same policy to tunneled node user traffic as you would to any untrusted wired traffic. The profile specified by the **aaa authentication wired** command determines the initial role, which contains the policy. The VLAN setting on the concentrator port must match the VLAN that will be used for users at the managed device.

Figure 21 *Tunneled Node Configuration Operation*



Configuration Procedures

The section below details different configuration procedures related to tunneled-node configuration:



ArubaOS does not allow a tunneled-node client and tunneled-node server to co-exist on the same managed device at the same time. The managed device must be configured as either a tunneled-node client or a tunneled-node server. By default, the managed device behaves as a tunneled-node server. However, once tunneled-node-server xxx.xxx.xxx.xxx is configured on the managed device, the managed device becomes a tunneled-node client. To remove the tunneled-node client function, use the command **tunneled-node-server 0.0.0.0** to disable the tunneled-node client on the managed device.

Configuring a Wired Tunneled Node Client

This section describes the CLI procedures to configure a tunneled node client.

1. Access the Wired tunneled node CLI according to the instructions provided in the installation guide that shipped with your tunneled node. Console access (9600 8N1) and SSH access are supported.
2. Specify the IP address of the managed device and specify tunnel loop prevention.

```
(host)(mynode)(config) #tunneled-node-address <tnode-ip-address>
(host)(mynode)(config) #tunnel-loop-prevention
```

3. Access each interface that you want to use and assign it as a tunneled node port.
```
(host)[mynode](config) #interface gigabitethernet <slot/module/port>
(host)[mynode](config-submode) #tunneled-node-port
```
4. Verify the configuration.
```
(host)[mynode](config-submode) #exit
(host)[mynode](config) #show tunneled-node config
```

Configuring an Access Port as a Tunneled Node Port

You can configure any port on any managed device as a tunneled node port using the **tunneled-node-port** command. Set the tunneled-node-address as the managed device to act as the tunneled node termination point. The **tunneled-node-port** command tells the physical interface to tunnel that traffic on the managed device.

1. Enable portfast on the wired tunneled node.
```
(host)(mynode)(config) #interface gigabitethernet <slot/module/port>
(host)(mynode)(config) #spanning-tree portfast
```
2. Assign a VLAN to the tunneled node port.
```
(host)[mynode](config-submode) #switchport mode access
(host)[mynode](config-submode) #switchport access vlan <id>
```

Configuring a Trunk Port as a Tunneled Node Port

- To enable switchport on the wired tunneled node execute the following commands:
```
(host)[mynode](config-submode) #switchport mode trunk
(host)[mynode](config-submode) #switchport trunk allowed vlan <WORD>
```
- To verify the status of the wired tunneled node execute the following commands:
(host)[mynode](config-submode) #show tunneled-node state
```
(host)[mynode](config-submode) #show tunneled-node config
```
- To check the current usage on the managed device execute the following command:
```
(host)[mynode] #show license-usage ap
```

  Each tunneled-node client uses one AP license. Attaching an additional wired client on the tunneled node client does not increment the AP license usage on the managed device.

Dynamic Segmentation

The Dynamic Segmentation solution is Aruba's ability to assign policy (roles), to a wired port based on the access method of a client. Further, using ClearPass Policy Manager, we can add context such as time-of-day and type-of-machine. The solution also provides users the ability to segment client traffic via traditional, locally-switched VLANs or to tunnel traffic back to an Aruba Mobility Controller. The two types of tunneling that are present in Dynamic Segmentation are:

- User-based tunneling
- Port-based tunneling

In the earlier releases this solution was called per user tunneled-node, which was built on top of Aruba's per-port tunneled node or port-based tunneling. Port-based tunneling allows the switch to tunnel traffic to an Aruba Mobility Controller on a per-port basis i.e., all traffic on a configured switch port was statically tunneled to an Aruba Mobility Controller. User-based tunneling in Dynamic Segmentation implements the capability to tunnel traffic on a user role-based or device basis, tunneling traffic of a given client or device based on an assigned user role. The policies associated with that client could be driven through a RADIUS server such as ClearPass Policy Manager, a downloaded role from ClearPass

Policy Manager, or by local MAC authentication in the switch. User-based tunneling can authenticate these devices using ClearPass Policy Manager, and tunnel the client traffic, utilizing the advanced firewall and policy capabilities in the Aruba Mobility Controller. It can also provide high availability and load balancing with controller clustering in ArubaOS 8.x.0.0, providing secure access to IoT devices within the Aruba Intelligent Edge wired network.



- User-based tunneled switch can only be terminated using the controller switch-ip. Non-controller switch-ip address termination is not supported including VRRP IP termination.
- The port-based tunneling feature is supported only if the native VLAN is configured on the switch trunk port.
- Overriding port-based tunnel client VLAN at controller is supported **only** for untagged VLANs configured on the port-based tunneling switch port and **is not supported** when both untagged and tagged VLAN is configured on the port-based tunneling switch port.

Starting from this release, IPv6 support is available for the Aruba Dynamic Segmentation solution.



Authentication is supported only from the switch, and not from the controller.

Jumbo MTU support is added to multicast tunnels to ensure that multicast stream from a multicast server to a Dynamic Segmentation user is not dropped. This change is effective for customers using ArubaOS 8.1.x.x or later versions and Aruba access switch version 16.08.

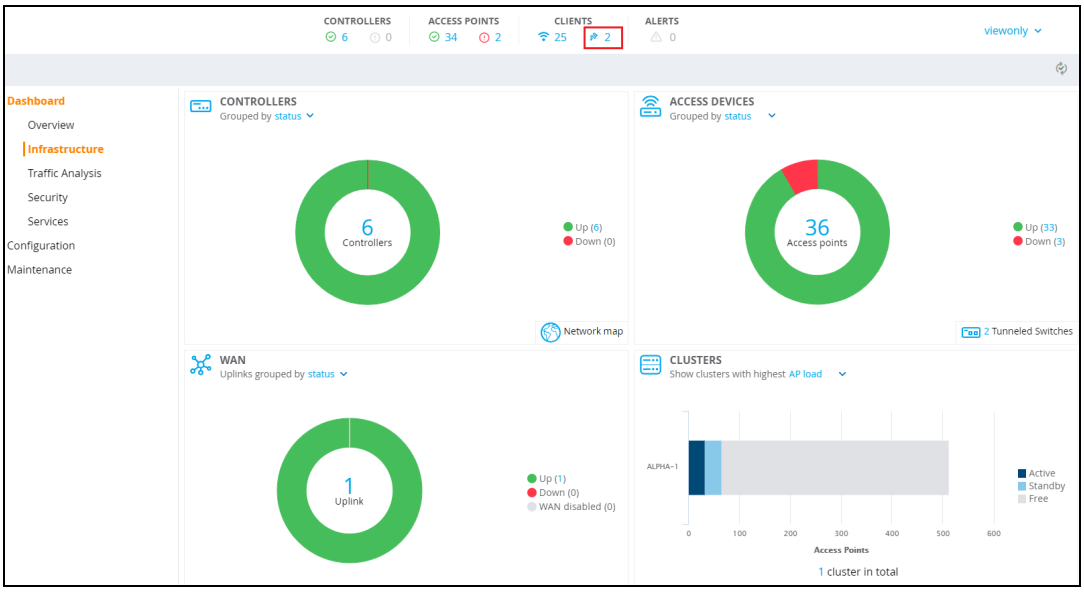


The Cluster Live Upgrade feature is not supported as part of the Dynamic Segmentation solution.

Dynamic Segmentation Visibility

Starting ArubaOS 8.4.0.0, wired clients are displayed in the WebUI providing visibility into the various aspects of the Dynamic Segmentation solution. Navigate to **Dashboard** under **Managed Network** node hierarchy. The number of wired clients connected are displayed at the top of the page under **CLIENTS** label. Clicking the link displays a page that will display details of the client.

Figure 22 *Wired Clients Visibility*

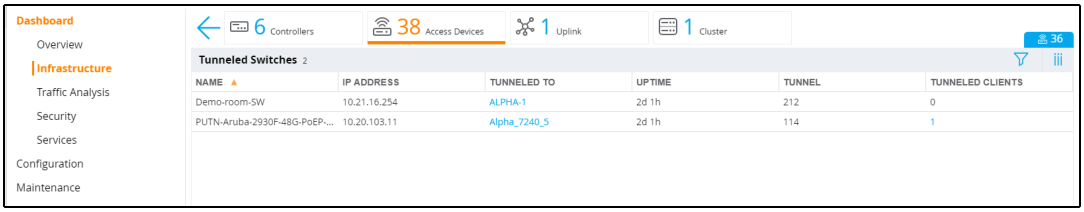


Tunneled Switches Table

Starting ArubaOS 8.4.0.0, tunneled switches table is introduced to display information on the tunneled switches.

- When an Aruba access switch is connected to a managed device that is part of a cluster, clicking a value in the **TUNNELED TO** column will redirect the user to **Dashboard > Infrastructure > Clusters** page.
- When an Aruba access switch is connected to a managed device that is part not of a cluster, clicking a value in the **TUNNELED TO** column will redirect the user to **Dashboard > Infrastructure > Controllers** page.
- Clicking an Active or Standby managed device will redirect the user to **Dashboard > Infrastructure > Controllers** page.
- Clicking a value in the **TUNNELED TO** column will redirect the user **Dashboard > Overview > Clients** page. The **Customize columns** filter enables the user to apply appropriate filters.

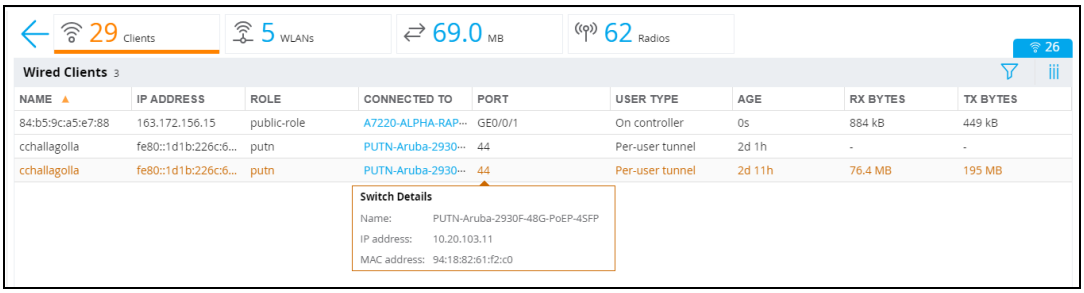
Figure 23 *Tunneled Switches Table*



Wired Clients Table

Starting ArubaOS 8.4.0.0, user-based tunneled node users are displayed in the wired clients table. Hovering the mouse over the port of a switch displays **Name**, **IP address**, and **MAC address** of the switch. Navigate to **Dashboard** under **Managed Network** node hierarchy and select the wired clients displayed at the top of the page under **CLIENTS** label. The wired clients table is displayed.

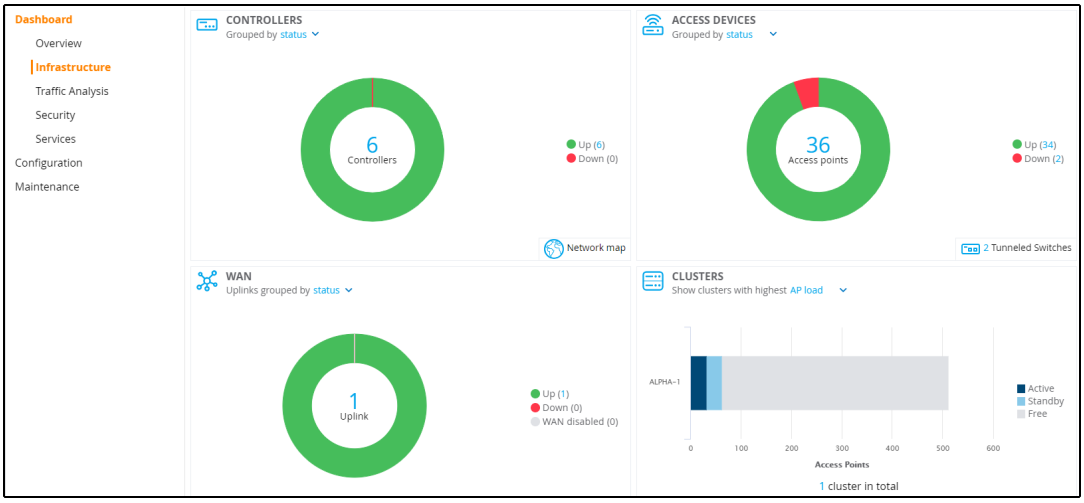
Figure 24 *Wired Clients Table*



Tunneled Switches Displayed on AP Page

Starting ArubaOS 8.4.0.0, when at least 1 tunneled switch is connected to a managed device within the selected network node, the icon and number of **Tunneled Switches** is displayed in the bottom right corner of the **Access Points** card.

Figure 25 *Tunneled Switches*



Deployment and Usage of VLANs in Dynamic Segmentation

In the current deployments, a VLAN is required to be configured on the ArubaOS switch and also on the controller. This also means that the same VLAN must be configured across the network and used for the same purposes across the deployment.

In large enterprise deployments it is not possible to have all the VLANs across all the switches and controllers across all the buildings carrying the same subnets. Starting ArubaOS 8.4.0.0, VLAN configuration is no longer required on ArubaOS switches for Dynamic Segmentation users. User VLANs is no longer configured on ArubaOS switch, but role-based VLANs will be used to assign roles for Dynamic Segmentation users so that user's traffic gets classified in some VLAN.

Supported Deployments

Stand-alone controller or managed devices supports user-based tunnel node feature in stand-alone and cluster environments to allow certain class of wired users to have their traffic inspected through DPI and firewall. The class of wired users whose traffic need to be redirected will be determined by the authenticating server, ClearPass Policy Manager.

User authentication will be performed by the ArubaOS switch where the ClearPass Policy Manager server returns:

- Only user role and switch has configuration of 'redirect secondary role' for the ClearPass Policy Manager returned role.
- User role along with redirect attribute with secondary role.

On successful authentication, the ClearPass Policy Manager server provides a **redirect** attribute with secondary user role to ArubaOS switch. This allows the switch to tunnel the user traffic to the stand-alone controller or managed devices for further inspection and these wired users are known as **Tunneled Users** on the stand-alone controllers or managed devices.



Ensure to configure user-based tunneled node defined VLAN and user-role on the managed device else the user will be rejected.

This feature is enabled by default on the managed device. This feature can be enabled or disabled on the ArubaOS switch. To enable this feature, refer to *HP's Management Configuration Guide 16.04*.

You can use the following show command to check if the user-based tunnel node is configured and is working as expected:

```
(host) [mynode] #show tunneled-node-mgr
  cluster-bucket-map      Information on Cluster Bucket Map
  cluster-node-list       Information on Cluster Node List
  gsm-counters            Information on GSM counters
  node-heartbeat-table    Information on Node Heartbeat Tables
  stats                   Information on Tunneled Node Manager Statistics
  trace-buf               Shows contents of trace buffer
  tunnel-vlan-user-map    Information on User count on each Tunnel Vlan pair
  tunneled-nodes          Information on Tunneled Nodes
  tunneled-users          Information on Tunneled Users
  user-tunnel-table       Information on User Tunnel Tables
```



Dynamic Segmentation functionality will not be enabled if you are running ArubaOS-Switch version 16.08 connected to a Mobility Controller Virtual Appliance or cluster running ArubaOS 8.3.0.0. The Mobility Controller Virtual Appliance must be upgraded to ArubaOS 8.4.0.0.

Support for Downloadable User Roles in Cluster Deployments

Starting from ArubaOS 8.4.0.0, support is extended for downloadable user roles in cluster deployments. This feature provides a seamless redundancy for dynamic policy assignments.

In this deployment, each ArubaOS switch establishes a connection with the active managed device and a secondary connection to the stand-by managed device. This allows the applied client role to be automatically replicated on the secondary managed device, thus minimizing the risk of clients loosing connectivity if the active managed device gets disconnected.

In case of downloadable user role implementation for WLAN clients, APs will maintain two connections with the cluster and not with ArubaOS switches.

ClearPass Policy Manager will be used to define the roles and policies, which will be downloaded to the managed devices in the cluster that is performing Dynamic Segmentation. ArubaOS switches that have Dynamic Segmentation activated on the port will also have downloadable role support.



If multiple authentication servers that are configured in a managed device point to the same CPPM address, which is used for CPPM role download, then all authentication servers should be updated with the CPPM credentials.

Support for Downloadable Roles for User-Based Tunneled Node Users

Starting from ArubaOS 8.3.0.0, this feature allows the managed device to get the user role from the Aruba ClearPass Policy Manager server while tunneling wired user's traffic to the managed device. That is, the ClearPass Policy Manager downloadable role feature is integrated with user-based tunneled node users.

When the user is successfully authenticated, ClearPass Policy Manager server sends two VSA attributes to the ArubaOS switch. First VSA contains the ClearPass Policy Manager primary user role to be applied on the switch and the second VSA contains redirect attribute with secondary user role. The ArubaOS switch sends the secondary user role name to the managed device and once this information is provided by the ArubaOS switch, the managed device starts the user role download process and after a successful download, the managed device applies the user role policies.

For more information on configuration of ClearPass Policy Manager authentication server, see Configuring Username and Password for ClearPass Policy Manager Authentication.

Support for IGMP and MLD Proxies for User-Based Tunneled Node Users

Starting from ArubaOS 8.1.0.1, IGMP and MLD proxies are supported for user-based tunneled node users in a cluster setup.



The cluster-profile multicast VLAN configuration is not supported for the user-based tunneled node feature. If it is configured, user-based tunneled node users cannot receive multicast stream.

Licensing Requirements

Starting with ArubaOS 8.4.0.0, user-based tunneled users will need a license for Dynamic Segmentation to function. The Aruba access switch will be viewed as an AP from the managed device's perspective. The user needs to procure a license for each Aruba access switch similar to procuring a license for an AP. If the license is not installed, Aruba access switch will not be allowed to form tunnels to Mobility Controllers running ArubaOS 8.4.0.0.



Licenses are consumed per switch not per port. A switch stack is considered a single switch for licensing purposes.

The following are deployment scenarios to consider:

- **10 APs and 10 switches**: Customer is using only AP licenses. For the deployment to function, 20 AP licenses are needed.
- **10 APs and 10 switches**: Customer is using AP, PEF, and RF Protect licenses. For deployment to function, 20 AP licenses, 20 PEF licenses, and 20 RF Protect licenses are needed.

The ArubaOS software allows you to use an external authentication server or the internal user database of the Mobility Master to authenticate clients who need to access the wireless network.

The following sections provide a general overview of the Mobility Master authentication server management:

- Understanding Authentication Server Best Practices and Exceptions on page 192
- Understanding Servers and Server Groups on page 192

Configuring Authentication Servers and Server Groups

The following topics describe the procedures to create and manage external and internal authentication servers and server groups.

- Configuring Authentication Servers on page 193
- Managing the Internal Database on page 207
- Configuring Server Groups on page 208
- Assigning Server Groups on page 214
- Configuring Authentication Timers on page 219
- Authentication Server Load Balancing on page 220
- Testing a Configured Authentication Server on page 220

Understanding Authentication Server Best Practices and Exceptions

For an external authentication server to process requests from Mobility Master, you must configure the server to recognize the Mobility Master. Refer to the vendor documentation for information on configuring the authentication server.

To configure Microsoft IAS and Active Directory, see the following links:

- http://technet2.microsoft.com/windowsserver/en/technologies/ias.mspx
- http://www.microsoft.com/en-us/server-cloud/windows-server/active-directory.aspx

Understanding Servers and Server Groups

Mobility Master supports the following external authentication servers:

- RADIUS
- LDAP
- TACACS+
- Windows (For stateful NTLM authentication)

NOTE

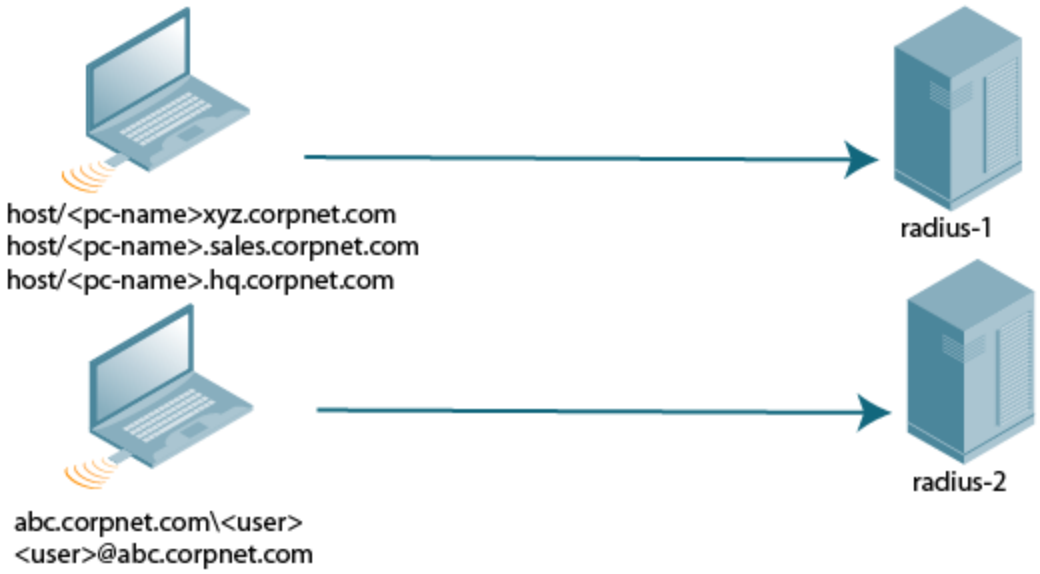
A maximum of 128 LDAP, RADIUS, and TACACS servers, each can be configured on a managed device.

Additionally, you can use the internal database to authenticate users by creating entries for users, their passwords, and their default role.

You can create *groups* of servers for specific types of authentication. For example, you can specify one or more RADIUS servers to be used for 802.1X authentication. The list of servers in a server group is an ordered list. This means that the first server in the list is always used unless it is unavailable, in which case the next server in the list is used. You can configure servers of different types in one group. For example, you can include the internal database as a backup to a RADIUS server.

Figure 26 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 26 *Server Group*



Server names are unique. You can configure the same server in multiple server groups. You must configure the server before you can add it to a server group.

NOTE

If you use the internal database for user authentication, use the predefined “Internal” server group.

You can also include conditions for server-derived user roles or VLANs in the server group configuration. The server derivation rules apply to all servers in the group.

Configuring Authentication Servers

This section describes how to configure RADIUS, LDAP, TACACS+ and Windows external authentication servers and the internal database.

This section includes the following information:

- Configuring a RADIUS Server on page 194
- RADIUS Service-Type Attribute on page 196
- Enabling Radsec on RADIUS Servers on page 196

- Configuring Username and Password for ClearPass Policy Manager Authentication on page 202
- Configuring an RFC-3576 RADIUS Server on page 202
- Configuring an LDAP Server on page 204
- Configuring a TACACS+ Server on page 205
- Configuring a Windows Server on page 206

Configuring a RADIUS Server

The following procedure describes how to configure a RADIUS server:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** tab.
2. In the **All Servers** table, click **+** to add a new server.
3. Enter a **Name** and **IP address/ hostname** for the new server and set the **Type** to **RADIUS**. Click **Submit**.
4. In the **All Servers** table, select the server created to configure server parameters.
5. Enter the parameters as described in Table 32. Select the **Mode** check box to activate the authentication server.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.



The RADIUS configuration on the Mobility Master under the **/mm** node is used only for management authentication on the Mobility Master and not for user or device (wired or wireless) authentication. The configuration under the **/mm** node is pushed only to the redundant Mobility Master pair and not to managed devices. Configuring RADIUS servers for clients or managed devices should be done on or under the **/md** node.

The following CLI commands configure a RADIUS server:

```
(host) [mynode] (config) #aaa authentication-server radius <name>
   host <ipaddr>
   key <psk>
   enable
```

Table 32: *RADIUS Server Configuration Parameters*

Parameter	Description
Name	Name of the RADIUS server.
IP address/ hostname	IP address or FQDN of the authentication server. The maximum supported FQDN length is 63 characters. Default: N/A
Auth Port	Authentication port of this server. Default: 1812
Acct Port	Accounting port of this server. Default: 1813
Shared key	Shared secret between the managed device and the authentication server. The maximum length is 128 characters.

Parameter	Description
	Default: N/A
Retype key	Re-enter the shared key.
Timeout	Maximum time, in seconds, that the managed device waits before timing out the request and resending it. Default: 5 seconds
Retransmits	Maximum number of retries sent to the server by the managed device before the server is marked as down. Default: 3
NAS ID	NAS identifier to use in RADIUS packets.
NAS IP	The NAS IP address to be sent in RADIUS packets from that server. **NOTE:** If you define a local NAS IP using the **Configuration > Security > Authentication > Servers** page and also define a global NAS IP using the **Configuration > Security > Authentication > Advanced** page, the global NAS IP address takes precedence.
Enable IPv6	Enable or disable IPv6 for this server. Default: Disabled
NAS IPv6	The NAS IPv6 address to be sent in RADIUS packets.
Use MD5	Use MD5 hash of cleartext password. Default: Disabled
Mode	Enables or disables the server. Default: Enabled
Lowercase MAC addresses	Send MAC address with lowercase in the authentication and accounting requests to this server. Default: Disabled
Use IP address for calling station ID	Enables or disables using the IP address for the calling station ID. Default: Disabled
MAC address delimiter	Send MAC address with the following delimiters in the authentication and accounting requests of this server: ▪ **colon**: Send MAC address as XX:XX:XX:XX:XX:XX ▪ **dash**: Send MAC address as XX-XX-XX-XX-XX-XX ▪ **none**: Send MAC address as XXXXXXXXXXXX ▪ **oui-nic**: Send MAC address as XXXXXX-XXXXXX Default: none
Service-type of FRAMED-USER	Send the service-type as FRAMED-USER instead of LOGIN-USER. For more information, see RADIUS Service-Type Attribute on page 196. Default: Disabled
CPPM credentials	Use the ClearPass Policy Manager server authentication.

Parameter	Description
CPPM username	Enter the ClearPass Policy Manager username.
CPPM password	Enter the ClearPass Policy Manager password.
Retype password	Re-enter the ClearPass Policy Manager password.
Station ID type	Select one of the following the station ID types to be sent with the RADIUS attribute, Called Station ID for authentication and accounting requests.: ▪ **AP group** ▪ **AP MAC address** ▪ **AP name** ▪ **IP address** ▪ **MAC address** ▪ **VLAN ID**
Station ID delimiter	Select one of the following delimiter options: ▪ **Colon** ▪ **Dash** ▪ **None**
Include SSID	Select the check box to include the SSID name in the Called Station ID attribute.

.

RADIUS Service-Type Attribute

Managed devices send the following Service-Type attribute values for RADIUS authentication requests.

Table 33: *RADIUS Service-Type Attributes*

RADIUS Attribute	Authentication Type	Attribute Value
Service-Type	MAC	Call-Check
	802.1X	Framed
	Captive Portal	Login

The service-type-framed-user configuration of the RADIUS server overwrites all the attribute values to Framed irrespective of the authentication type. Existing deployments that depend upon this attribute for their third-party RADIUS integrations should make changes to support these new service types.

Enabling Radsec on RADIUS Servers

Conventional RADIUS protocol offers limited security. This level of limited security is not sufficient for authentication that takes place across unsecured networks such as the Internet. To address this, the RADIUS over TLS or Radsec enhancement is introduced to ensure RADIUS authentication and accounting data is transmitted safely and reliably across insecure networks. The default destination port for RADIUS over TLS is TCP/2083. Separate ports are not used for authentication, accounting, and dynamic authorization changes.

In a TLS connection, both the managed device (TLS client) and the Radsec server (TLS server) need to authenticate each other using certificates. For the managed device to authenticate the Radsec server:

- The CA certificate should be uploaded as a **Trusted CA** if the Radsec server uses a certificate signed by a CA.
- Self-signed certificates should be uploaded as a **PublicCert** if the Radsec server uses a self-signed certificate.



If neither of these certificates are configured, the managed device does not try to establish any connection with the Radsec server, even if Radsec is enabled.

The managed device must also send a TLS client certificate to the Radsec server by uploading a certificate on Mobility Master as **ServerCert** and configuring Radsec to accept and use the certificate. If a certificate is not configured, Mobility Master uses the device certificate in its TPM. In this case, the Aruba device CA that signed the certificate should be configured as a Trusted CA on the Radsec server.



When Radsec support is enabled, the default RADIUS shared key is **radsec** and remains the same even if the user configures a different shared key.

The following CLI commands configure Radsec on RADIUS server:

```
(host) [mynode] (config) #aaa authentication-server radius <rad_server_name>
  enable-radsec
  radsec-client-cert-name <name>
  radsec-port <radsec-port>
  radsec-trusted-cacert-name <radsec-trusted-ca>
  radsec-trusted-servercert-name <name>
```

To upload certificates through the CLI, see Managing Certificates on page 972.



To configure a Radsec server as RFC 3576 server for dynamic CoA, see Enabling Radsec on RADIUS Servers on page 196.

RADIUS Server VSAs

VSAs are a method for communicating vendor-specific information between Network Access Servers and RADIUS servers, allowing vendors to support their own extended attributes. You can use Aruba VSAs to derive the user role and VLAN for RADIUS-authenticated clients; however the VSAs must be present on your RADIUS server. This requires that you update the RADIUS dictionary file with the vendor name (Aruba) and/or the vendor-specific code (14823), the vendor-assigned attribute number, and the attribute format (such as string or integer) for each VSA. For more information on VSA-derived user roles, see Workflow for Assigning a User Role on page 504



Starting from ArubaOS 8.4.0.0, the RADIUS server VSAs support Aruba-Captive-Portal-VSA attribute.

For the current and complete list of all RADIUS VSAs available in the version of ArubaOS currently running on your Mobility Master, access the command-line interface and issue the command **show aaa radius-attributes**.

Bandwidth-VSAs

Starting from ArubaOS 8.2.0.0, the managed device can dynamically assign per-user or per-group bandwidth rate on Layer 3 authenticated clients based on the direction from RADIUS server. To direct the managed device to enforce bandwidth rates for specific clients after successful Captive-Portal

authentication, three RADIUS Vendor-Specific Attributes named Bandwidth-VSAs are added in the RADIUS Access-Accept packet.

Table 34: *Bandwidth-VSAs*

VSA	Type	Value	Description
Nomadix-Group-Bw-Policy-ID	Integer	19	Set to zero for per-client, else the group-ID for per-group.
WISPr-Bandwidth-Max-Up	Integer	7	Upstream bandwidth rate in bits per second.
WISPr-Bandwidth-Max-Down	Integer	8	Downstream bandwidth rate in bits per second.
Vendor ID	Integer	8	ID of the vendor.



The server-redirected bandwidth control feature supports only D-tunnel and controller wired clients.

The following CLI command checks the Dynamic Bandwidth Contracts currently assigned:

```
(host) # show aaa bandwidth-contracts dynamic
```

Customizing the RADIUS Attributes

Starting from ArubaOS 8.1.0.0, the users can now configure RADIUS modifier profile to customize the attributes that are included, excluded and modified in the RADIUS request before it is sent to the authentication server. The RADIUS modifier profile can be configured and applied to either Access-Request or Accounting-Request or both on a RADIUS authentication or accounting server.

This profile can contain up to 64 RADIUS attributes with static values that are used either to add or update in the request and another 64 RADIUS attributes to be excluded from the Requests.

Two new parameters have been added in the RADIUS modifier profile :

- **auth-modifier**: When assigned, it references to a RADIUS modifier profile which is applied to all Access-Requests sending to this RADIUS authentication server.
- **acct-modifier**: When assigned, it references to a RADIUS modifier profile which is applied to all Accounting-Requests sending to this RADIUS accounting server.

You can create a RADIUS modifier profile to customize the attributes that are included, excluded and modified in the RADIUS request before it is sent to the authentication or accounting server.

The following procedure describes how to create a RADIUS modifier profile and customize the RADIUS attributes:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Under **All Profiles**, expand **Wireless LAN**.
3. Click **Radius Modifier**.
4. Under the **Radius Modifier Profile: New Profile**, click **+** to add a Radius modifier profile.
5. Enter the **Profile name**.

6. In **+Attr** field, click **+** and select a name from **Name** drop-down list box and set the **Type** to Static and enter the **Static_val.**Click **OK**. The name field should be available in the list of attributes when we execute the command, **show aaa radius-attribute** command
7. In the **-Attr** field, click **+** and select the name of the attribute you want to exclude from **-attr** drop-down list box and click **OK**.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands create a RADIUS modifier profile and customize the RADIUS attributes:

```
(host) [md] (config) #aaa authentication-server radius radius1
(host) [md] (RADIUS Server "radius1) #
  acct-modifier
  acctport
  auth-modifier
  authport
  …
  …
(host) [md] (config) #aaa radius modifier <profile_name>
  clone
  exclude
  include
  no
(host) [md] #show aaa radius modifier <profile_name>
```

Dynamic Data Support

Starting from ArubaOS 8.2.0.0, support for dynamic data for the included attributes in the RADIUS Attribute modifier is supported. Users can configure the dynamic value for each included attribute in the RADIUS modifier to be one or two data items. Following data items can be picked to form the dynamic value for each included attribute:

- **AP-Name**: Name of the AP which the client currently associated to.
- **AP-MAC-Address**: MAC-address of the AP which the client currently associated to.
- **AP-Group**: Group-name of the AP which the client currently associated to.
- **ESSID**: ESSID which the client currently associated to.

Field1 and Field2 have the same value but these can be used for different combination with the delimiter. This included attribute are of type String and can contain up to 128 bytes.

The following procedure describes how to configure a RADIUS modifier profile with single-item dynamic data:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration** > **Systems** > **Profiles** tab.
2. Under **All Profiles**, expand **Wireless LAN**.
3. Click **Radius Modifier.**
4. In the **Radius Modifier Profile: New Profile**, click **+** to add a new radius modifier profile.
5. Enter a **Profile name**.
6. Click **+** in **+Attr** field and select a name from the **Name** drop-down list and set the **Type** to **dynamic.**
7. Select the first dynamic field from the **D_field1** drop-down list.
8. (Optional) Select the second dynamic field from the **D_field2** drop-down list.
9. Select the delimiter from the **D_delimiter** drop-down list.

10. . Click **OK.**
11. Click **Submit**
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure a RADIUS modifier profile with single-item dynamic data, :

```
(host)(config) #aaa radius modifier dynamic-mod
(host) (Radius Modifier Profile "dynamic-mod") #?
  clone           Copy data from another Radius Modifier Profile
  exclude         Attribute to be excluded in RADIUS request
  include         Attribute/Value to be included in RADIUS request
  no              Delete Command

(host) (Radius Modifier Profile "dynamic-mod") #include ?
  <name>          RADIUS Attribute Name

(host) (Radius Modifier Profile "dynamic-mod") #include Aruba-Location-Id ?
  dynamic         First dynamic field
  static          Static Data

(host) (Radius Modifier Profile "dynamic-mod") #include Aruba-Location-Id dynamic ?
  ap-group1       Use AP group as first dynamic field
  ap-macaddr1     Use AP mac address as first dynamic field
  ap-name1        Use AP name as first dynamic field
  essid1          Use essid as first dynamic field
  user-vlan1      Use user's current VLAN-ID as first dynamic field

(host) (Radius Modifier Profile "dynamic-mod") #include Aruba-Location-Id dynamic ap-
name1
```

To configure a RADIUS modifier profile with two-item dynamic data

```
(host) (Radius Modifier Profile "dynamic-mod") #include Aruba-Location-Id dynamic ?
  ap-group1       Use AP group as first dynamic field
  ap-macaddr1     Use AP mac address as first dynamic field
  ap-name1        Use AP name as first dynamic field
  essid1          Use essid as first dynamic field
  user-vlan1      Use user's current VLAN-ID as first dynamic field

(host) (Radius Modifier Profile "dynamic-mod") #include Aruba-Location-Id dynamic essid1
?
  with            Optional second dynamic field

(host) (Radius Modifier Profile "dynamic-mod") #include Aruba-Location-Id dynamic essid1
with ?
  ap-group2       Use AP group as second dynamic field
  ap-macaddr2     Use AP mac address as second dynamic field
  ap-name2        Use AP name as second dynamic field
  essid2          Use essid as second dynamic field
  user-vlan2      Use user's current VLAN-ID as first dynamic field

(host) (Radius Modifier Profile "dynamic-mod") #include Aruba-Location-Id dynamic essid1
with ap-macaddr2 ?
  delimiter       Delimiter between fields

(host) (Radius Modifier Profile "dynamic-mod") #include Aruba-Location-Id dynamic essid1
with ap-macaddr2 delimiter ?
  at              Use '@' as delimiter between fields
  colon           Use ':' as delimiter between fields
  dash            Use '-' as delimiter between fields
  dollar          Use '$' as delimiter between fields
  hash            Use '#' as delimiter between fields
```

```
   none             NULL
   percent          Use '%' as delimiter between fields
   semicolon        Use ';' as delimiter between fields
   slash            Use '/' as delimiter between fields
   space            Use ' ' as delimiter between fields

(host) (Radius Modifier Profile "dynamic-mod") #include Aruba-Location-Id dynamic essid1
with ap-macaddr2 delimiter at ?
```

The following CLI command shows a RADIUS modifier profile with mixing of static- and dynamic- data:

```
(host) (config) #show aaa radius modifier dynamic-mod
Radius Modifier Profile
-----------------------
Action Attribute Name   Data Type Data Value
------ --------------   --------- -----------
+Attr Aruba-Location-Id  dynamic  essid1 with ap-macaddr2 delimiter at
+Attr BW-Area-Code       static   "212"
+Attr BW-City-Name       static   "San Jose"
+Attr Aruba-AP-Group     dynamic  ap-group1
-Attr Aruba-Device-Type
```

Dynamically Assign VLAN-ID to NAS-Port

The following CLI command configures a RADIUS modifier to assign the client's VLAN-ID to the NAS-Port RADIUS attribute:

```
(host) [mode] (config) # aaa radius modifier "Hilton-Eleven"
   include "NAS-Port-ID" dynamic user-vlan1
!
```

The following CLI command assigns the RADIUS modifier to a RADIUS authentication server:

```
(host) [mode] (config) #aaa authentication-server radius "eleven-server"
   .....
   auth-modifier "Hilton-Eleven”
   .....
!
```

RADIUS Server Authentication Codes

A configured RADIUS server returns the following standard response codes.

Table 35: *RADIUS Authentication Response Codes*

Code	Description
0	Authentication OK.
1	Authentication failed : user/password combination not correct.
2	Authentication request timed out : No response from server.
3	Internal authentication error.
4	Bad Response from RADIUS server : verify shared secret is correct.
5	No RADIUS authentication server is configured.
6	Challenge from server (This does not necessarily indicate an error condition).

RADIUS Server Fully Qualified Domain Names

If you define a RADIUS server using the FQDN of the server rather than its IP address, the managed device periodically generates a DNS request and caches the IP address returned in the DNS response.

To view the IP address that currently correlates to each RADIUS server FQDN, access the command-line interface in config mode and issue the **show aaa fqdn-server-names** command.

DNS Query Intervals

If you define a RADIUS server using the FQDN of the server rather than its IP address, the managed device periodically generates a DNS request and caches the IP address returned in the DNS response. DNS requests are sent every 15 minutes by default.

You can use either the WebUI or the CLI to configure how often a DNS request is generated to cache the IP address for a RADIUS server identified via its FQDN.

The following procedure describes how to configure DNS query intervals:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Advanced** page.
2. Expand the **DNS Query Interval** accordion, enter a new DNS query interval from 1-1440 minutes, in the **DNS Query Interval** (min) field.
3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures DNS query intervals:

```
(host) [mynode] (config) #aaa dns-query-interval <minutes>
```

Configuring Username and Password for ClearPass Policy Manager Authentication

Authentication to ClearPass Policy Manager is enhanced to use configurable usernames and passwords instead of a support password. The support password is vulnerable to attacks as the server certificate presented by ClearPass Policy Manager server is not validated.

Configuring an RFC-3576 RADIUS Server

You can configure a RADIUS server to send user disconnect, CoA, and session timeout messages as described in RFC 3576, “Dynamic Authorization Extensions to Remote Dial In User Service (RADIUS).”



For Remote AP, RADIUS CoA is supported on tunnel and split-tunnel forwarding modes only.

For Campus AP, RADIUS CoA is supported on tunnel and decrypt-tunnel forwarding modes only.

The disconnect, session timeout, and CoA messages sent from the server to a managed device contains information to identify the user for which the message is sent. Starting from ArubaOS 8.5.0.0, the managed device also accepts disconnect, session timeout, and CoA requests from IPv6 address based DAC, and identifies user sessions based on the user's IPv6 address. Mobility Master supports the following attributes for identifying the users who authenticate with an RFC 3576 server:

- **user-name**: name of the user to be authenticated
- **framed-ip-address**: user IPv4 address
- **framed-ipv6-address**: user IPv6 address
- **calling-station-id**: phone number of a station that originated a call
- **accounting-session-id**: unique accounting ID for the user session.



The IPv4 address has a higher priority over IPv6 address for identification of user sessions.

If the authentication server sends both supported and unsupported attributes to a managed device, the unknown or unsupported attributes are ignored. If no matching user is found, a *503: Session Not Found* error message is sent back to the RFC 3576 server.

The following procedure describes how to configure the RFC-3576 RADIUS server:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** tab.
2. To define a new RFC 3576 RADIUS server, click **+** under **All Servers**.
3. Select the **Type** as **Dynamic Authorization** from the drop-down list.
4. In **IP address version**, select either **IPv4** or **IPv6** radio button based on your preference.
5. In **IP address**, enter the IPv4 or IPv6 address.
6. Click **Submit**.
7. From the **All Servers** list, select the server that you created to configure the server parameters.
8. Under **Server Options**, enter the server authentication key into the **Key** and **Retype key** fields.
9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the RFC-3576 RADIUS server:

```
(host) [mynode] (config) #aaa rfc-3576-server <ipaddr>
  clone <source>
  key <psk>
  no ...
```

Configuring an RFC 3576 RADIUS Server with Radsec

Starting from ArubaOS 8.2.0.0 new enhancements in RFC 3576 will make the communication between disconnect requests from the RADIUS server and managed device comprehensive. This release also supports detection of duplicate disconnect requests from the RADIUS server. This change ensures that:

1. The system is less prone to a packet replay attack.
2. The managed device does not process duplicate disconnect requests from the RADIUS server multiple times.

Starting from ArubaOS 8.5.0.0, you can also use IPv6 address to configure RFC 3576 RADIUS server with Radsec.

The following CLI commands configure an RFC 3576 RADIUS Server with Radsec:

```
(host) [mynode] (config) #aaa rfc-3576-server <ipaddr>
  clone <source>
  enable-radsec
  event-timestamp-requi..
  key <psk>
  no ...
  replay-protection
  window-duration
  enable-radsec
  no ...
```

The following enhancements have been introduced in RFC 5176:

State Attributes

State attributes are part of CoA request and not disconnect requests. Disconnect requests from the RADIUS server have many optional attributes and service type attributes is one such attribute. If the value of this attribute is “authorize only”, RFC 5176 has made it mandatory to have a state attribute in the disconnect requests and if state attribute is missing, **Error cause 402** is reported.

Error Cause 407

RFC 5176 has introduced this new error cause for disconnect requests responses. **407 - Notify DAC** for the invalid attribute value associated to any attribute.

Duplicate Request Detection

Disconnect requests detect duplicate requests coming from the same RADIUS server IP address or source port, for the packet with same sequence number. The minimum time span between two disconnect requests from same source can be configured and any two requests within this time window is considered duplicate, which is rejected.

Service Type Attribute: Authorize Only

If Network Access Service receives a disconnect request from Dynamic Authorization Client, with the service-type attribute **Authorize Only**, then Network Access Service should send Dynamic Authorization Client-Negative acknowledgment to the Dynamic Authorization Client, since service-type attribute can only be a part of CoA requests and not disconnect requests.

Configuring an LDAP Server

The following table describes the parameters that you configure for an LDAP server.

Table 36: *LDAP Server Configuration Parameters*

Parameter	Description
Host	IP address of the LDAP server. Default: N/A
Admin-dn	Distinguished name for the admin user who has read/search privileges across all the entries in the LDAP database (the user does need write privileges, but will be able to search the database, and read attributes of other users in the database).
Admin-passwd	Password for the admin user. Default: NAAN
Re-type admin-passwd	Re-enter the admin password.
Allow Clear-Text	Allows clear-text (unencrypted) communication with the LDAP server. Default: disabled
Auth port	Port number used for authentication. Default: 389
Base-dn	Distinguished Name of the node that contains the entire user database. Default: N/A
Filter	A string searches for users in the LDAP database. The default filter string is: **(objectclass=*)**. Default: N/A

Parameter	Description
Key Attribute	A string searches for a LDAP server. For Active Directory, the value is sAMAccountName. Default: sAMAccountName
Timeout	Timeout period of a LDAP request, in seconds. Default: 20 seconds
Mode	Enables or disables the server. Default: enabled
Preferred Connection Type	Preferred type of connection between a managed device and the LDAP server. The default order of connection type is: 1. clear-text 2. ldap-s 3. start-tls The managed device first attempts to contact the LDAP server using the preferred connection type, and only attempts to use a lower-priority connection type if the first attempt is not successful. **NOTE:** If you select **clear-text** as the preferred connection type, you must also enable the **allow-cleartext** option.
Maximum number of non-admin connections	Configure the maximum number of non-admin connections to the server. Default: 4
Chase referral	Chase referrals anonymously.

The following procedure describes how to configure an LDAP server:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** tab.
2. To configure an LDAP server, click **+** under **All Servers**. Enter the **Name** and **IP address / hostname** of the server. Set the **Type** to **Ldap** and click **Submit**.
3. Select the name of the server created to configure server parameters. Enter parameters as described in Table 36. Select the **Mode** check box to activate the authentication server.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.



The configuration does not take effect until you perform this step.

The following CLI command configures an LDAP server:

```
(host) [mynode] (config) #aaa authentication-server ldap <name>
   host <ipaddr>
```

(enter parameters as described in Table 36)

```
   enable
```

Configuring a TACACS+ Server

Table 37 defines the TACACS+ server parameters.

Table 37: *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
Re-type Key	Re-enter the key.
TCP Port	TCP port used by server. Default: 49
Retransmits	Maximum number of times a request is retried. Default: 3
Timeout	Timeout period for TACACS+ requests, in seconds. Default: 20 seconds
Mode	Enables or disables the server. Default: enabled
Session Authorization	Enables or disables session authorization. Session authorization turns on the optional authorization session for admin users. Default: disabled

The following procedure describes how to configure a TACACS+ server:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** tab.
2. To configure a TACACS+ server, click **+** under **All Servers**. Enter the **Name** and **IP address** of the server. Set the **Type** to **TACACS** and click **Submit**.
3. Select the server created to configure server parameters. Enter the parameters as described in Table 37. Select the **Mode** check box to activate the authentication server.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.



The configuration does not take effect until you perform this step.

The following CLI commands configure a TACACS+ server and session authorization:

```
(host) [mynode] (config) #aaa authentication-server tacacs <name>
   clone default
   host <ipaddr>
   key <psk>
   enable
   session-authorization
```

Configuring a Windows Server

The following table defines parameters for a Windows server used for stateful NTLM authentication.

Table 38: *Windows Server Configuration Parameters*

Parameter	Description
Host	IP address of the server. Default: N/A
Mode	Enables or disables the server. Default: enabled
Windows Domain	Name of the Windows Domain assigned to the server.

The following procedure describes how to configure a Windows server:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** tab.
2. To configure a Windows server, click **+** under **All Servers**. Enter the **Name** and **IP address / hostname** of the server. Set the **Type** to **Windows** and click **Submit**.
3. Select the server created to configure server parameters. Enter the parameters as described in Table 38.
4. Select the **Mode** check box to activate the authentication server.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.



The configuration does not take effect until you perform this step.

The following CLI commands configure a Windows server:

```
(host) [mynode] (config) #aaa authentication-server windows <windows-server-name>
  host <ipaddr>
  enable
```

Managing the Internal Database

You can create entries in the internal database to authenticate clients. The internal database contains a list of clients, along with the password and default role for each client. When you configure the internal database as an authentication server, client information is checked in incoming authentication requests against the internal database.

Configuring the Internal Database

Mobility Master uses the internal database for authentication by default. You can choose to use the internal database in a managed device by entering the CLI command **aaa authentication-server internal use-local-switch**. If you use the internal database in a managed device, you need to add clients on the managed device.

The following CLI command configures the internal database:

```
(host) [mynode] #local-userdb add {generate-username|username <name>}{
generate-password|password <password>}
```

Managing Internal Database Files

Mobility Master allows you to import and export user information tables to and from the internal database. These files should not be edited once they are exported. Mobility Master only supports the

importing of database files that were created during the export process. Note that importing a file into the internal database overwrites and removes all existing entries.

The following CLI commands configure the import and export of internal database files:

```
(host) [mynode] #local-userdb export <filename>
(host) [mynode] #local-userdb import <filename>
```

Configuring Server Groups

You can create *groups* of servers for specific types of authentication. For example, you can specify one or more RADIUS servers to be used for 802.1X authentication. You can configure servers of different types in one group. For example, you can include the internal database as a backup to a RADIUS server. You can also configure the same server in more than one server group. However, you must configure the server before you can include it in a server group using the WebUI or the CLI.

The following procedure describes how to configure a server group:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** tab.
2. The **Server Groups** table displays the server group list.
3. Click **+** in the **Server Groups**. Enter the name of the new server group and click **Submit**.
4. Select the new server group created.
5. In **Server Group <server group name>**, click the **Servers** tab and click **+** to add a server to the group.
   a. To add an existing server, select **Add existing server** and choose a server from the list. Click **Submit**.
   b. To add a new server, select **Add new server**. Specify a server type from the **Type** drop-down list, and enter a **Name** and **IP address/hostname** for the server. Click **Submit**.
   c. Repeat the above step(s) to add other servers to the group.
6. Click **Submit**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.
8. Click **Pending Changes**.

The following CLI command configures a server group:

```
(host) [mynode] (config) #aaa server-group <name>
  auth-server <name>
```

Configuring Server List Order and Fail-Through

The servers in a server group are part of an ordered list. The first server in the list is always used by default, unless it is unavailable, in which case the next server in the list is used. You can configure the order of servers in the server group through the WebUI using the **up** or **down** arrows (the top server is the first server in the list). In the CLI, the **position** parameter specifies 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 also enable *fail-through* authentication for the server group so that if the first server in the list returns an authentication deny, the managed device attempts authentication with the next server in the ordered list. The managed device attempts to authenticate with each server in the list until 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 a managed device (AAA FastConnect).
- Enabling this feature for a large server group list may cause excess processing load on the managed device. It is recommended that you use server selection based on domain matching whenever possible (see Configuring Dynamic Server Selection on page 210).
- Certain servers, such as the RSA RADIUS server, lock out the managed device 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 containing a subset of the usernames and passwords used in the network. When you enable fail-through authentication, users that fail authentication with the first server on the list will be authenticated with the second server.

The following procedure describes how to configure the server list order and fail-through:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** tab.
2. Click **+**. Enter **ldap-1** for the **Name** of the server, enter the **IP address / hostname** for the server, and set the **Type** to **LDAP** and click **Submit**.
3. Click **+**. Enter **ldap-2** for the **Name** of the server, enter the **IP address / hostname** for the server, and set the **Type** to **LDAP** and click **Submit**.
4. Under **All Servers**, select **ldap-1** to configure server parameters. Select the **Mode** check box to activate the authentication server.
5. Click **Submit**.
6. Repeat step on page 209 to configure **ldap-2**.
7. Click **+** under the **Server Groups** table to add a new server group. Set the server group name to **corp-serv**, and then click **Submit**.
8. Select **corp-serv** from the **Server Groups** table to configure the server group settings.
9. In **Server group <corp-serv>**, select the **Options** tab.
10. Select the **Fail through** check box.
11. Click **Submit**.
12. Navigate to the **Servers** tab.
13. Click **+** to add a server to the group.
    a. Select **ldap-1**, and then click **Submit**.
    b. Repeat the step above to add **ldap-2** to the server group.
14. Click **Submit**.
15. Click **Pending Changes**.
16. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the server list order and fail-through:

```
(host)[mynode] (config) #aaa authentication-server ldap ldap-1
  host 10.1.1.234
(host) [mynode] (config) #aaa authentication-server ldap ldap-2
  host 10.2.2.234
```

```
(host) [mynode] (config) #aaa server-group corp-serv
   auth-server ldap-1 position 1
   auth-server ldap-2 position 2
   allow-fail-through
```

Configuring Dynamic Server Selection

Managed devices 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)

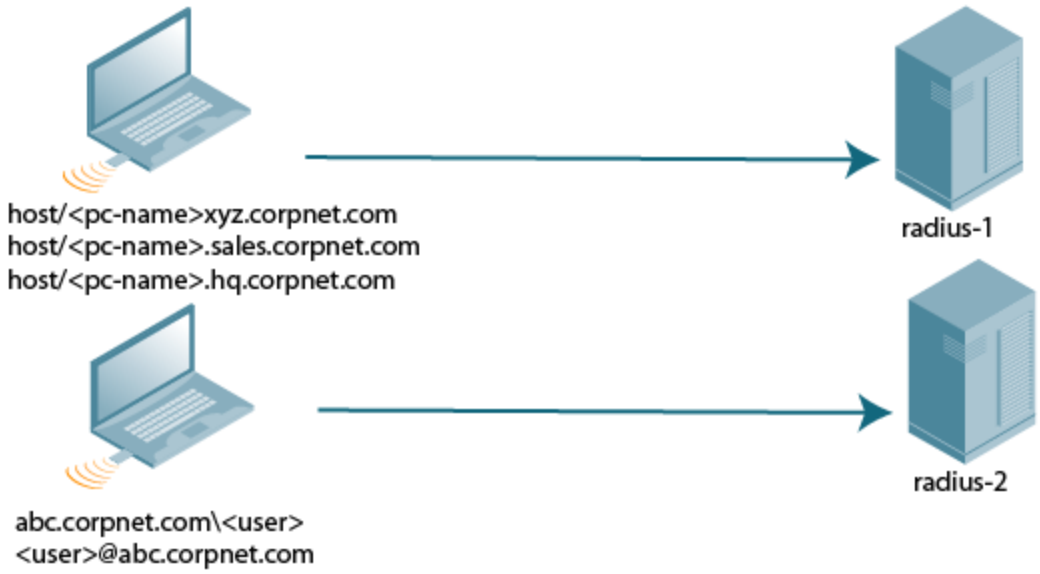
When you configure a server in a server group, you have the option to associate the server with one or more match rules. A match rule for a server can be one of the following:

- 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. Managed devices compare 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 managed device 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.

Figure 27 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 27 *Domain-Based Server Selection Example*



The following procedure describes how to configure dynamic server selection:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** tab.
2. Select a server group from the **Server Groups** table.
3. In the **Server Group <server group name>,** select the **Server Rules** tab, click **+**.
   a. Select an attribute from the **Attribute** drop-down list.
   b. Select an **Operation** to apply a condition to the attribute.
   c. Set the **Operand** value to the client or user information.
   d. Set the **Action** to apply an action to the attribute.
   e. Set the **Role** to set a role to the attribute.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands configure dynamic server selection:

```
(host) [mynode] (config) #aaa server-group <group>
  auth-server <name> [match-authstring contains|equals|starts-with <string>] [match-fqdn
  <string>] [position <number>] [trim-fqdn]
```

Configuring Match FQDN Option

You can also use the "match FQDN (domain name)" option for a server rule. With this rule, the server is selected if the <domain> portion of the user information in the formats **<domain>\<user>** or **<user>@<domain>** matches a specified string *exactly*. 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.

The following procedure describes how to configure a match FQDN option:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** page.
2. Select a server group from the **Server Groups** table.
3. In the **Server Group <server group name>,** select the **Server Rules** tab and click **+**.
   a. Select **Domain-Name** from the **Attribute** drop-down list.
   b. Set the **Operation** to **equals**.
   c. Set the **Operand** value to the client or user information.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI command configures a match FQDN option:

```
(host) [mynode] (config) #aaa server-group <group>
  auth-server <name> match-fqdn <string>
```

Trimming Domain Information from Requests

Before a managed device 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 managed device in the following formats:

- **<domain>\<user>** : the <domain>\ portion is truncated
- **<user>@<domain>** : the @<domain> portion is truncated



This option does not support client information sent in the format host/<pc-name>.<domain>.

The following procedure describes how to configure the trimming domain information from requests:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** tab.
2. Select a server group from the **Server Groups** table.
3. Under **Server Group <server group name>,** click the **Servers** tab and select a server or click **+** to add a new server to the group.
   a. To add an existing server, select **Add existing server** and choose a server from the list. Click **Submit**.
   b. To add a new server, select **Add new server**. Specify a server type from the **Type** drop-down list, and enter a **Name and IP address / hostname** for the server. Click **Submit**.
4. Select the new server.
5. In **Server group <server group name> < server name>**, click the **Server Group Trim FQDN** tab**.**
6. Select the **Trim FQDN** check box.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI command configures the trimming domain information from requests:

```
(host) [mynode] (config) #aaa server-group <group>
   auth-server <name> trim-fqdn
```

Configuring Server-Derivation Rules

When you configure a server group, you can set the VLAN or role for clients based on attributes returned for the client by the server during authentication. The server derivation rules apply to all servers in the group. The user role or VLAN assigned through server derivation rules takes precedence over the default role and VLAN configured for the authentication method.



The authentication servers must be configured to return the attributes for the clients during authentication. For instructions on configuring the authentication attributes in a Windows environment using IAS, refer to the documentation 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 39 describes the server rule parameters you can configure.

Table 39: *Server Rule Configuration Parameters*

Parameter	Description
Attribute	This is the attribute returned by the authentication server that is examined for *Operation* and *Operand* match.
Operation	This is the match method by which the string in *Operand* is matched with the attribute value returned by the authentication server. ■ contains : The rule is applied if and only if the attribute value contains the string in parameter *Operand.* ■ starts-with : The rule is applied if and only if the attribute value returned starts with the string in parameter *Operand.* ■ ends-with : The rule is applied if and only if the attribute value returned ends with the string in parameter *Operand.* ■ equals : The rule is applied if and only if the attribute value returned equals the string in parameter *Operand.* ■ not-equals : The rule is applied if and only if the attribute value returned is not equal to the string in parameter *Operand.* ■ value-of : This is a special condition. What this implies is that the role or VLAN is set to the value of the attribute returned. For this to be successful, the role and the VLAN ID returned as the value of the attribute selected must already be configured on the managed device when the rule is applied.
Operand	This is the string to which the value of the returned attribute is matched.
Action	Defines whether to assign a role or a VLAN to the user when the rule is matched.
Role or VLAN	The server derivation rules apply to either user role or VLAN assignment. With Role assignment, a client can be assigned a specific role based on the attributes returned. In VLAN assignment, the client can be placed in a specific VLAN based on the attributes returned.

The following procedure describes how to configure Server-Derivation Rules:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** tab.
2. Select a server group from the **Server Groups** table .
3. In **Server Group <server group name >,** select the **Servers** tab and select a server or click **+** to add a new server to the group.
   a. To add an existing server, select **Add existing server** and choose a server from the list. Click **Submit**.
   b. To add a new server, select **Add new server**. Specify a server type from the **Type** drop-down list, and enter a **Name and IP address / hostname** for the server. Click **Submit**.

4. In the **Server Rules** tab, click **+** to add server derivation rules for assigning a user role or VLAN.
   a. Select the **Attribute** from the drop-down list.
   b. Select the **Operation** from the drop-down list.
   c. Enter the **Operand**.
   d. To set a role, select **set role** from the **Action** drop-down list. Select the role to be assigned from the **Role** drop-down list.
   e. To set a vlan, select **set vlan** from the **Action** drop-down list. Select the VLAN name or ID from the **Vlan** drop-down list.
   f. Click **Submit**.
   g. Repeat the above steps to add other rules for the server group.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands configure Server-Derivation Rules:

```
(host) [mynode] (config) #aaa server-group <name>
(host) [mynode] (Server Group name) #set {role|vlan} condition <attribute> contains|ends-
with|equals|not-equals|starts-with <operand> set-value <set-value-str> position <number>
```

Configuring a Role Derivation Rule for the Internal Database

When you add a user entry to the internal database, you can specify a user role (see Managing the Internal Database on page 207). 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:

The following CLI command configures a server derivation rule for the internal database:

```
(host) [mynode] (config) #aaa server-group internal
  set role condition Role value-of
```

Assigning Server Groups

You can create server groups for the following purposes:

- User authentication
- Management authentication
- Accounting

You can configure all types of servers for user and management authentication (see Table 40). Accounting is only supported with RADIUS and TACACS+ servers when RADIUS or TACACS+ is used for authentication.

Table 40: *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

The following section describes user authentication, management authentication, and accounting:

User Authentication

For information about assigning a server group for user authentication, refer to the *Roles and Policies* chapter of the *ArubaOS User Guide*.

Management Authentication

Users who need to access Mobility Master to monitor, manage, or configure the Aruba user-centric network can be authenticated with RADIUS, TACACS+, or LDAP servers or the internal database.



Only user record attributes are returned upon successful authentication. Therefore, to derive a management role other than the default mgmt auth role, set the server derivation rule based on the user attributes.

The following CLI command enables management authentication:

```
(host)[mynode] (config) #aaa authentication mgmt
   server-group <group>
   enable
```

Accounting

You can configure accounting for RADIUS and TACACS+ server groups.



RADIUS or TACACS+ accounting is only supported when RADIUS or TACACS+ is used for authentication.

The following section describes RADIUS accounting, roaming RADIUS accounting service, RADIUS accounting on multiple servers and TACACS+ accounting:

RADIUS Accounting

RADIUS accounting allows user activity and statistics to be reported from managed devices to RADIUS servers:

1. The managed device 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 acknowledgment of the packet.
2. The managed device sends an Accounting Stop packet when a user logs off; the packet information includes various statistics such as elapsed time, input and output bytes, and packets. The RADIUS server sends an acknowledgment of the packet.

The following attributes can be sent to a RADIUS accounting server:

- **Acct-Status-Type:** This attribute marks the beginning or end of accounting record for a user. Current values are Start, Stop, and Interim Update.
- **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 managed device. This is only sent in Accounting-Request records, where the Acct-Status-Type is Stop or Interim Update.
- **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 Mobility Master.
- **NAS-Identifier:** This is set in the RADIUS server configuration.
  - In the **Mobility Master** node hierarchy of the WebUI, navigate to the **Configuration > Authentication > Advanced** page. Under **RADIUS Client**, enter the IPv4 or IPv6 address.
- **NAS-IP-Address:** IP address of the managed device. You can configure a "global" NAS IP address:
  - 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 managed device.
- **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 managed device.

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 statistical attributes are sent only in Interim-Update and Accounting Stop packets (they are not sent in Accounting Start packets):

- Acct-Input-Octets
- Acct-Output-Octets
- Acct-Input-Packets
- Acct-Output-Packets
- Acct-Input-Gigawords
- Acct-Output-Gigawords

Remote APs in split-tunnel mode now support RADIUS accounting. If you enable RADIUS accounting in a split-tunnel Remote APs AAA profile, the managed device sends a RADIUS accounting start record to the RADIUS server when a user associates with the remote AP, and sends a stop record when the user logs out or is deleted from the user database. If interim accounting is enabled, the managed device sends updates at regular intervals. Each interim record includes cumulative user statistics, including received bytes and packets counters.

The following procedure describes how to assign a server group for RADIUS accounting:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > AAA Profiles** tab.
2. Expand the **AAA Profiles** pane and select the **default** profile instance.
3. (Optional) In the **AAA Profile: default** pane, select **RADIUS Interim Accounting** to allow the managed device to send Interim-Update messages with current user statistics to the server at regular intervals. This option is disabled by default, allowing the managed device to send only *start* and *stop* messages RADIUS accounting server.
4. Select a AAA profile, and then scroll down to select the **RADIUS Accounting Server Group** for the AAA profile. Select the **Server group** from the drop-down list.

You can add additional servers to the group or configure server rules.

5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands configure a server group for RADIUS accounting:

```
(host) [mynode] (config) #aaa profile <profile>
   radius-accounting <group>
   radius-interim-accounting
```

Roaming RADIUS Accounting Service

Starting from ArubaOS 8.1, the Roaming RADIUS Accounting Service creates an Accounting session for each wireless client. The records in the session contain the same set of RADIUS attributes as compared to the timer-based RADIUS Interim-Update Accounting record, except the statistics attributes. Whenever a wireless client roams to a different AP, the Roaming triggered RADIUS Interim-Update Accounting record is sent to the configured RADIUS Accounting server. This record is used to track the current location of the wireless client. Currently this feature is supported for wireless clients in both cluster and non-cluster environments, but is not supported for wired, VPN/VIA, and L3-Mobility clients.

The following procedure describes how to enable roaming RADIUS accounting services:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > AAA Profiles** tab.
2. Expand **AAA Profiles** and select a AAA profile instance.
3. In the **AAA Profile: <name of the profile>** pane, select the **RADIUS Roaming Accounting** check box.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI command enables roaming RADIUS accounting services:

```
(host) [mynode] (config) # aaa profile <profile_name>
   radius-accounting <group>
   radius-roam-accounting
```

The following CLI command checks if roaming-triggered RADIUS accounting is enabled:

```
(host) #show aaa profile <profile_name>
```

Configuring RADIUS Accounting on Multiple Servers

ArubaOS provides support to send RADIUS accounting to multiple RADIUS servers. Mobility Master notifies all the RADIUS servers to track the status of authenticated users. Accounting messages are sent to all the servers configured in the server group in a sequential order.

The following procedure describes how to enable multiple server account functionality:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > AAA Profiles** tab.
2. Expand **AAA Profiles** and select a AAA profile instance.
3. In the **AAA Profile: <name of the profile>** pane, select the **Multiple Server Accounting** check box.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI command enables RADIUS accounting on multiple servers functionality:

```
(host) [mynode] (config) # aaa profile <profile_name>
   multiple-server-accounting
```

TACACS+ Accounting

TACACS+ accounting allows commands issued on a Mobility Master or managed device to be reported to TACACS+ servers. You can specify which types of commands are reported (action, configuration, or show commands), or report all commands.

You can only configure TACACS+ accounting using the CLI.

The following CLI commands configure TACACS+ accounting:

```
(host) [mm] (config) #aaa tacacs-accounting
(host) ^[mm] (config-submode) #command {action|all|configuration|show}
(host) ^[mm] (config-submode) #server-group <name of the TACACS server>
(host) ^[mm] (config-submode) #write memory
```

Configuring Authentication Timers

The following procedure describes how to configure authentication timers:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Advanced** tab.
2. Expand **Authentication Timers** .
3. Configure the timers as described in Table 41.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Table 41: *Authentication Timers*

Timer	Description
User Idle Timeout	Maximum period after which a client is considered idle if there is no wireless traffic from the client. The timeout period is reset if there is wireless traffic. If there is no wireless traffic in the timeout period, the client is aged out. Once the timeout period has expired, the user is removed. If the keyword **seconds** is not specified, the value defaults to minutes at the command line. Range: 1–255 minutes (30–15300 seconds) Default: 5 minutes (300 seconds)
Authentication Server dead Time	Maximum period, in minutes, that the managed device considers an unresponsive authentication server to be “out of service.” This timer is only applicable if there are two or more authentication servers configured on a managed device. If there is only one authentication server configured, the server is never considered out of service, and all requests are sent to the server. If one or more backup servers are configured and a server is unresponsive, it is marked as out of service for the dead time; subsequent requests are sent to the next server on the priority list for the duration of the dead time. If the server is responsive after the dead time has elapsed, it can take over servicing requests from a lower-priority server; if the server continues to be unresponsive, it is marked as down for the dead time. Range: 0–50 minutes Default: 10 minutes
Logon User Lifetime	Maximum time, in minutes, unauthenticated clients are allowed to remain logged on. Range: 0–255 minutes Default: 5 minutes
User Interim stats frequency	Sets the timeout value for user stats, reporting in minutes or seconds. Range: 300-3600 seconds, or 5-60 minutes Default: 600 seconds

The following CLI commands configure timers that you can apply to clients. If the optional seconds keyword is not specified for the **idle-timeout** and **stats-timeout** parameters, the value defaults to minutes:

```
(host)[mynode] (config) #aaa timers
   dead-time <minutes>
   idle-timeout <time> [seconds]
   logon-lifetime <0-255>
   stats-timeout <time> [seconds]
```

Authentication Server Load Balancing

Load balancing of authentication servers ensures that the authentication load is split across multiple authentication servers, thus avoiding any one particular authentication server from being overloaded. Authentication Server Load Balancing functionality enables Mobility Master to perform load balancing of authentication requests destined for external authentication servers (RADIUS or LDAP). This prevents any one authentication server from having to handle the full load during heavy authentication periods, such as at the start of the business day.

Previously, the controller used the first authentication server in the server group list. The remaining servers in that group would be used in sequential order only when an authentication server was down. Thus, the controllers performed fail-over instead of load balancing of authentication servers.

The load balancing algorithm computes the expected time taken to authenticate a new client for each authentication server and chooses that authentication server with the shortest expected authentication time. The load balancing algorithm maintains re-authentication stickiness, meaning that at the time of re-authentication, the request is forwarded to the same server where it was originally authenticated.

Enabling Authentication Server Load Balancing Functionality

The following CLI command enables authentication server load balancing functionality:

```
(host) [mynode] (config) #aaa server-group <group>
   load-balance
   auth-server s1
   auth-server s2
```

The following CLI command disables load balancing:

```
(host) [mynode] (config) #aaa server-group <group>
   no load-balance
```



If you configure an internal server in the server group, load balancing is not applicable to the internal server. The Internal server will be used as a fall-back when all other servers in the group are down.

Testing a Configured Authentication Server

You can test the configured RADIUS authentication server in the WebUI or the CLI. This feature allows you to check a configured RADIUS authentication server or the internal database. You can use this feature to check for an “out of service” RADIUS server.

The following CLI commands configure a user in the internal database:

```
(host)(mynode)# local-userdb add kgreen lkjHGfds
(host)(mynode)# aaa test-server pap internal kgreen lkjHGfds
```

Starting from ArubaOS 8.1.0.0, the **aaa test-server** command includes the verbose option. The **verbose** option helps display the response of the RADIUS server on a successful or failed authentication. This eases troubleshooting an active network. This enhancement applies to both the WebUI and the CLI.

The following procedure describes how to get the RADIUS server responses on an authentication success or failure:

1. In the **Mobility Master** node hierarchy, go to the **Diagnostics > Tools > AAA Server Test** tab.
2. Select a server from the **Server Name** drop-down list.
3. Select an option for **Authentication method**. You can select either **MSCHAPv2** or **PAP**.
4. Enter the user credentials in the **Username** and **Password** text boxes.
5. Click **Test**. The Authentication Status along with the RADIUS server response is displayed.

The following CLI command displays the RADIUS server attributes as returned by the server:

```
(host)(mynode) # aaa test-server mschapv2 internal raduser1 raduser verbose
Authentication Successful
Processing time (ms) : 1.397
Attribute value pairs in response
---------------------------------
Vendor  Attribute  Value
------  ---------  -----
        MS-CHAPv2
        Role       guest
```

This chapter describes how to configure MAC-based authentication on the Mobility Master using the WebUI or the CLI.

Use MAC-based authentication to authenticate devices based on their physical MAC address. Although this not the most secure and scalable method, MAC-based authentication implicitly provides an additional layer of security to authenticate 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 through station A, then one method of authenticating station A is MAC-based. Clients may be required to authenticate themselves using other methods depending on the network privileges required.

MAC-based authentication can also be used to authenticate Wi-Fi phones as an additional layer of security to prevent other devices from accessing the voice network using what is normally an insecure SSID.

This chapter describes the following topics:

- Configuring MAC-Based Authentication on page 222
- Configuring Clients on page 223

Configuring MAC-Based Authentication

Before configuring MAC-based authentication, you must configure the following options:

- User role—The user role that will be assigned as the default role for the MAC-based authenticated clients. (See Roles and Policies on page 491 for information on firewall policies to configure roles.)
- Configure the default user role for MAC-based authentication in the AAA profile. If derivation rules exist or if the client configuration in the internal database has a role assigned, these values take precedence over the default user role.
- Authentication server group—The authentication server group that the managed device uses to validate the clients. The internal database can be used to configure the clients for MAC-based authentication. See Configuring Clients on page 223 for information on configuring the clients on the local database. For information on configuring authentication servers and server groups, see Authentication Servers on page 192.

The following section describes how to configure the MAC authentication profile:

Configuring the MAC Authentication Profile

The following procedure describes how to configure MAC-based authentication:

1. In the **Mobility Master** node hierarchy, select a managed device.
2. Navigate to the **Configuration > Authentication > L2 Authentication** tab.
3. Click **MAC Authentication**.
4. In the **MAC Authentication Profile: New Profile** window, click **+** to create a new profile.
5. Enter a **Profile name**.

6. Configure the parameters, as described in Table 42.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following table describes the parameters that you can configure for MAC-based authentication.

Table 42: *MAC Authentication Profile Configuration Parameters*

Parameter	Description
Profile name	Name of the MAC authentication profile.
Delimiter	Delimiter used in the MAC string: ■ colon specifies the format XX:XX:XX:XX:XX:XX ■ dash specifies the format XX-XX-XX-XX-XX-XX ■ none specifies the format XXXXXXXXXXXX ■ oui-nic specifies the format XXXXXX-XXXXXX Default: none
Case	The case (upper or lower) used in the MAC string. Default: lower
Max Authentication failures	Number of times a station can fail to authenticate before it is blacklisted. A value of zero disables blacklisting. Default: zero (0)
Reauthentication	Select the **Reauthentication** check box if you want to enable Reauthentication; Default: disable.
Reauthentication Interval	Time duration between reauthentication attempts. Configure a value in the range of 60–86,400. Reauthentication timer is configured in terms of seconds.
Use Server provided Reauthentication Interval	Select the **Use Server provided Reauthentication Interval** check box to use the interval provided by the server; Default: disable.

The following CLI commands configure a MAC authentication profile from the Mobility Master node:

```
(host)[mynode](config) #aaa authentication mac <profile>
(host) [mynode] (MAC Authentication Profile "profile") #case {lower|upper}
(host) [mynode] (MAC Authentication Profile "profile") #clone {default|<source>}
(host) [mynode] (MAC Authentication Profile "profile") #delimiter
{colon|dash|none|oui-nic}
(host) [mynode] (MAC Authentication Profile "profile") #max-authentication-failures
<max-authentication-failures-number>
(host) [mynode] (MAC Authentication Profile "profile") #reauthentication
(host) [mynode] (MAC Authentication Profile "profile") #timer reauth-period <reauth
period>
```

Configuring Clients

You can create entries in the Mobility Master's internal database to authenticate client MAC addresses. The internal database contains a list of clients along with the password and default role for each client. To configure entries in the internal database for MAC authentication, enter the username and password for each client.

The following section describes how to configure clients entries in the internal database.

The following procedure describes how to configure the clients:

1. In the **Mobility Master** node, navigate to the **Configuration > Authentication > Auth Servers** tab.
2. In the **All Servers** section, select **Internal**. The **Server > Internal** section is displayed below the **All Servers** section.
3. In the **Server > Internal > Users** tab, click **+** to open the **Internal Server > Add New User** section.
4. Click **Generate** beside the **User name** and **Password** text boxes, for automatic username and password generation. Otherwise, enter the username and password in the text boxes.
5. Select the role from the **Role** drop-down list.
6. Enter the email address in the **E-mail** text box.
7. Select the **Enabled** check box to activate the user entry on creation.
8. Select the expiration duration mode from the **Expiration** drop-down list. Expiration represents the maximum time duration that a guest account is valid for.
   - If you selected duration, set the time for expiration in minutes.
   - If you selected time, set the date (mm/dd/yyyy format) and time (hh:mm format) in the **Date** and **Time** boxes.
9. Click **Submit** at the bottom of the page.
10. Click **Pending Changes** at the top of the page.
11. In the **Pending Changes** window, select the check box indicating the pending change and click **Deploy Changes**.

The following CLI command configures the clients from the Mobility Master node:

```
(host)[mynode] #local-userdb add generate-username generate-password  role <user-role>
mode {disable} expiry {duration <1+> | time <mm/dd/yyyy> <hh:mm>}
```

Scaling numbers for Mobility Controller Platforms

When a client (wired or wireless) terminates on a managed device, the MAC address of the client is added to the user table. An IP user is created for the client when it obtains the IP address and communicates with the managed device. The maximum number of IP users allowed for each client or each MAC address (including IPv4 and IPv6 addresses) is 150% of the maximum MAC user for a given platform.



Ensure that the scaling number for IP users does not exceed 150% of the maximum MAC user for each platform. The managed device stops accepting IP users if the scaling number for IP users goes beyond 150% of the maximum MAC users.

Multi Pre-Shared Key

Multi Pre-Shared Key (MPSK) is an enhancement to WPA2-Personal that allows device-specific and group-specific passphrases. This offers enhanced security and deployment flexibility for headless and IoT devices over traditional per-SSID, static passphrases.

Passphrases can be administratively assigned to groups of devices in policy based on common attributes like device profiling data or uniquely assigned to each device during ClearPass Policy Manager device registration.

MPSK requires the WPA2-PSK operating mode with AES encryption. MAC Authentication via RADIUS or RadSec is used for passphrase authorization and role assignment between ClearPass Policy Manager and the managed device.



MPSK requires ClearPass Policy Manager v6.8.

For information on how to configure a SSID profile, see Configure the SSID profile for the configuration node.

Many distributed enterprises with branch and remote offices and locations use cost-effective hybrid WAN connectivity solutions that include low-cost DSL, 4G and LTE technologies, rather than relying solely on traditional E1/T1 or T3/E3 dedicated circuits. 7000 Series Cloud Services Controllers are optimized for these types of locations, which are more likely to use cloud security architectures instead of dedicated security appliances, and where clients are likely to access applications in the cloud, rather than on local application servers.

This chapter describes ArubaOS features designed to optimize the configuration and performance of managed devices in branch and remote offices, and lists the procedures to configure these features.

Learn more about Managed Device Optimization

Select any of the links below to view detailed information about ArubaOS features for managed device configuration and management, and examples of deployment topologies that support these features.

- Managed Device Feature Overview on page 226
- Zero-Touch Provisioning Overview on page 228
- WAN Authentication Survivability Overview on page 232
- Managed Device WAN Dashboard on page 236

Provision and Configure a Managed Device

The following sections describe the procedures to configure your network for zero-touch managed device provisioning, and to define configuration settings for a group of managed devices.

- Using ZTP with DHCP to Provision a Managed Device on page 242
- Health Check Services for Managed Devices on page 244
- WAN Optimization Through IP Payload Compression on page 245
- WAN Interface Bandwidth Priorities on page 246
- Uplink Monitoring and Load Balancing on page 247
- Hub and Spoke VPN Configuration on page 252
- IP Routes Configuration on page 254
- Uplink Routing using Next-hop Lists on page 255
- Policy Based Routing on page 257
- Address Pool Management on page 259
- Configuring WAN Authentication Survivability on page 263
- Preventing WAN Link Failure on Virtual APs on page 264
- Managed Device Integration with a Palo Alto Networks Portal on page 265

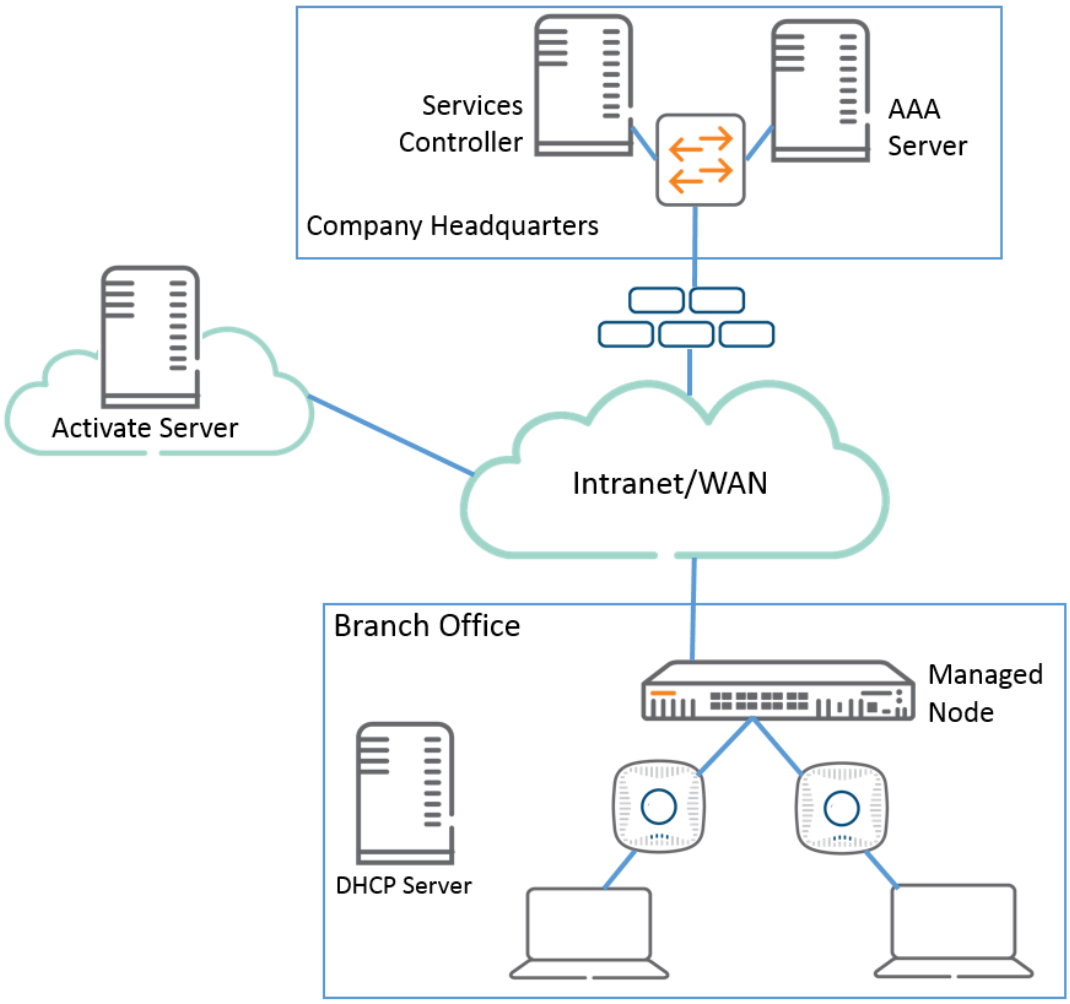
Managed Device Feature Overview

ArubaOS supports these distributed enterprises through the following features designed specifically for managed devices in branch and remote offices:

- Authentication survivability allows managed devices to store user access credentials and key reply attributes whenever clients are authenticated with external RADIUS servers or LDAP authentication servers, providing authentication and authorization survivability when remote authentication servers are not accessible.
- Integration with existing Palo Alto Networks Firewalls, like WildFire™ anti-virus and anti-malware detection services. In deployments with multiple Palo Alto Networks firewalls, managed devices can select the best PAN firewall based on priority and availability.
- Policy-based routing on each uplink interface, which allows you specify the next hop to which packets are routed. ArubaOS supports multiple next-hop lists, to ensure connectivity in the event that a device on the list becomes unreachable.
- Uplink and VPN redundancy, and per-interface bandwidth contracts to limit traffic for individual applications (or categories of applications) either sent from or received by a selected interface.
- Packet compression between Aruba devices (such as devices at the branch and main office), to maximize the amount of data that can be carried by the network.
- A WAN health-check feature that uses ping-probes to measure WAN availability and latency on each uplink.

The following diagram depicts a managed device topology where a managed device in the branch office learns the address, routing information, and other provisioning information from the Mobility Master.

Figure 28 *Managed Device Topology*



Scalable Site-to-Site VPN Tunnels

ArubaOS supports site-to-site IPsec tunnels based on an FQDN. When you identify the remote peer for a managed device using an FQDN, that node configuration can be applied across multiple branch managed devices, as the configured FQDN can resolve to different IP addresses for each local branch, based on local DNS settings.

Crypto maps for site-to-site VPNs support a VLAN ID as the identifier for the source network. When the VPN settings are pushed to a managed device, the IKE negotiation process uses the IP address range for the VLAN. This feature allows multiple managed devices to use a single group of configuration settings defined at a configuration node, as each managed device negotiates a different source network IP for its VLAN, based on the IP pool for the managed devices defined for that configuration node.

WAN Health Check

The health-check feature uses ping-probes to measure WAN availability and latency on selected uplinks. Based upon the results of this health-check information, the managed device can continue to use its primary uplink, or failover to a backup link. Latency is calculated based on the round-trip time of ping responses. The results of this health check appear in the **WAN** section of the Monitoring Dashboard.

IPsec Tunnels using GCM ciphers

Starting from ArubaOS 8.6.0.0, an IPsec tunnel can be established between managed devices and APs using GCM ciphers. The IPsec tunnel can be established without loading the ECDSA custom certificates. By default, the APs send the GCM cipher algorithm in the IPsec set, along with the current cipher list. New dynamic maps are programmed on the managed devices to establish the IPsec tunnels with GCM ciphers.

To establish a successful IPsec tunnel with GCM ciphers, disable the **default-rap-ipsecmap dynamic map** and ensure that there is an ACR license for each AP in the deployment.



220 Series and 550 Series access points do not support GCM ciphers. The IPsec tunnels are established using AES ciphers.

Zero-Touch Provisioning Overview

Traditionally, the deployment of controllers was a multiple step process where the master controller information and local configurations were first pre-provisioned. After the managed device connected to the network, it established a secure tunnel to the master and downloaded the global configuration. ZTP automates deployment of managed devices plug-n-play. The managed device now learns the required information from the network and provisions itself automatically. ArubaOS allows a managed device to automatically get its local and global configuration and license limits from Mobility Master.

This section includes the following topics:

- Why use ZTP?
- Managed Device Provisioning Modes
- Managed Device Address Pools
- Zero-Touch Provisioning Workflows
- ZTP Support Matrix



For more information about the procedures to prepare your network for ZTP, see Using ZTP to Provision a Managed Device.

Why use ZTP?

ZTP offers the following advantages over a standard managed device configuration:

- Simple deployment
- Reduced operational cost
- Limits to provisioning errors

A managed device configured using ZTP automatically discovers the Mobility Master, downloads its local configuration from that Mobility Master, and is provisioned with its device role, and country code.



The local configuration is the configuration that is specific to a managed device. That is, not the global configuration shared by a network of managed devices. This includes, but is not limited to, IP addresses and VLANs.

Once the managed device is provisioned, it is ready to obtain its global configuration in either of two ways:

- The administrator enters the global configuration via the WebUI or CLI of the Mobility Master.
- The managed device retrieves its global configuration from the Mobility Master.

Device-specific configurations that are common across multiple devices can be modified from a central location using the bulk edit feature. Users can apply common device configurations to a group of devices without having to update each device individually. Bulk edit supports, but is not limited to, the following configurations:

- Time zone
- Daylight savings time setting
- VLANs
- Managed device IP addresses
- DHCP pools

Managed Device Provisioning Modes

The administrator has the choice of provisioning modes that select how the managed device is supplied with its own IP address, role, country code, and configuration settings.

Once the managed device learns the IP address of the primary Mobility Master, the managed device contacts that Mobility Master and retrieves its configuration from its assigned configuration node.



Before you deploy a managed device, use you must create a configuration for that device at a configuration node on Mobility Master. Mobility Master pushes this configuration to the managed device when the device becomes active on the network.

ArubaOS supports the following provisioning modes for managed devices:

- **auto:** In this mode, the managed device:
  - obtains its IP address from DHCP
  - obtains its role, country code, and the IP addresses of the Mobility Master and any defined secondary Mobility Master from a provisioning rule in Activate
  - retrieves its configuration from a configuration node on Mobility Master

- **mini-setup:** In this mode, the managed device:
  - has its role set to local (local) when mini-setup is initiated
  - obtains its IP address from DHCP
  - is configured through the console with its country code and the IP address of the primary Mobility Master and (optionally) the secondary Mobility Master IP
- retrieves its local configuration group from the primary Mobility Master
- **full-setup:** In this mode, the managed device:
  - is configured with its role set to local (local) through the console
  - is configured to obtain its IP address through manual configuration of a static IP, DHCP, or PPPoE
  - is configured through the console with its country code and the IP address of the primary Mobility Master and (optionally) the secondary Mobility Master IP
  - retrieves its configuration from a configuration node on the primary Mobility Master

Managed Device Address Pools

Each managed device needs a pool of addresses it can dynamically assign to APs or users on each of its VLANs, and a separate IP address that managed device uses to create a GRE tunnel to Mobility Master. Mobility Master can assign IP these addresses to managed devices using dynamic address pools. These pools allow network administrators to create a generic configuration that provisions managed device interfaces with individual settings that are unique across branch offices. If managed devices are also serving as DHCP servers for other devices at that location, smaller DHCP pools for those individual branches can be dynamically carved out from a larger DHCP pool.

ArubaOS 8.0.0.0 supports three different types of address pools that can be applied to a hierarchy node

- **NAT Pools:** A NAT pool is used to assign IP addresses to a VLAN interface on a managed device . The range of addresses in this pool is available for use for any DHCP-enabled managed device when it is added to that specific node in the configuration hierarchy. When you add a managed device, a group of IP addresses is removed from the NAT pool on that hierarchy node and is and leased to the device. The IP addresses in a NAT pool are dynamic (leased) rather than static (permanently assigned), so addresses no longer in use are automatically returned to the pool for reallocation.
- **Tunnel pools:** A tunnel pool defines a range of IP addresses that can be used by the managed devices to create a GRE tunnel to the Mobility Master. When you add a managed device controller, an IP address is removed from the tunnel pool on that hierarchy node and is and leased to that device. Addresses no longer in use are automatically returned to the pool for reallocation.
- **VLAN pools:** A VLAN pool allocates a block of IP addresses for each managed device. The managed device acts as a DNS proxy server and dynamically assigns IP addresses from its allocated pool to each AP or client on the VLAN. A VLAN pool allocates multiple addresses to each managed device VLAN, unlike the tunnel pool, which assigns a single tunnel IP address to each managed device.

Zero-Touch Provisioning Workflows

The managed device obtains its IP address through DHCP by sending a DHCP discover on the default uplink port. The default uplink port is configured as an access port in VLAN 4094.

Next it will attempt to retrieve the provisioning parameters from Activate. If the managed device is unsuccessful in retrieving the provisioning parameters from Activate, it will retry in 30 seconds. The managed device keeps trying to retrieve the provisioning parameters from Activate every 30 seconds until it is successful or the administrator interrupts Auto-Provisioning by initiating mini-setup or full-setup.

To interrupt the auto provisioning process, enter the string **mini-setup** or **full-setup** at the initial setup dialog prompt shown below.

```
Auto-provisioning is in progress. Choose one of the following options to override or
debug...
   'enable-debug' : Enable auto-provisioning debug logs
   'disable-debug': Disable auto-provisioning debug logs
   'mini-setup' : Stop auto-provisioning and start mini setup dialog for smart-local role
   'full-setup' : Stop auto-provisioning and start full setup dialog for any role
Enter Option (partial string is acceptable):_
```

ZTP Support Matrix

The following table provides information about the ZTP support for the platforms.

Table 43: *ZTP Support for 7000 Series Platforms*

ArubaOS Version	7005	7008	7010	7030	7024
ArubaOS 8.2.2.6	0/0/3	0/0/7	0/0/15	0/0/7	0/0/23
ArubaOS 8.3.0.9	0/0/3	0/0/7	0/0/15	0/0/7	0/0/23
ArubaOS 8.4.0.4	All ports except 0/0/1	All ports except 0/0/1	All ports except 0/0/1	All ports except 0/0/1	All ports except 0/0/1
ArubaOS 8.5.0.2	0/0/3	0/0/7	All ports except 0/0/1	All ports except 0/0/1	All ports except 0/0/1
ArubaOS 8.6.0.0	All ports except 0/0/1	All ports except 0/0/1	All ports except 0/0/1	All ports except 0/0/1	All ports except 0/0/1
ArubaOS 8.7.0.0	All ports except 0/0/1	All ports except 0/0/1	All ports except 0/0/1	All ports except 0/0/1	All ports except 0/0/1

Table 44: *ZTP Support for 7200 Series and 9000 Series Platforms*

ArubaOS Version	7205	7210	7220	7240	7280	9004	9012
ArubaOS 8.2.2.6	0/0/0	0/0/1	0/0/5	0/0/5	N/A	N/A	N/A
ArubaOS 8.3.0.9	0/0/3	0/0/1	0/0/1	0/0/1	N/A	N/A	N/A
ArubaOS 8.4.0.4	0/0/3	All ports except 0/0/1	All ports except 0/0/1	All ports except 0/0/1	N/A	N/A	N/A
ArubaOS 8.5.0.2	0/0/3	All ports except 0/0/1	All ports except 0/0/1	All ports except 0/0/1	N/A	0/0/0	N/A
ArubaOS 8.6.0.0	0/0/0	All ports except 0/0/0	All ports except 0/0/1	All ports except 0/0/1	N/A	0/0/0	0/0/0

ArubaOS Version	7205	7210	7220	7240	7280	9004	9012
ArubaOS 8.7.0.0	0/0/0	All ports except 0/0/0	0/0/0	All ports except 0/0/1	N/A	0/0/0	0/0/0

WAN Authentication Survivability Overview

Authentication survivability is critical to managed device WLANs since most managed devices use geographically remote authentication servers to provide authentication and authorization services. When those authentication servers are not accessible, clients cannot access the WLAN because the managed device cannot authenticate them. ArubaOS authentication survivability allows managed devices to provide client authentication and authorization survivability when remote authentication servers are not accessible. When this feature is enabled, ArubaOS stores user access credentials and key reply attributes whenever clients are authenticated with external RADIUS servers or LDAP authentication servers. When external authentication servers are not accessible, the managed device uses its internal survival server to continue providing authentication and authorization functions by using the user access credentials and key reply attributes that were stored earlier.

When authentication survivability is enabled, an internal survival server on the managed node performs authentication functions, as well as EAP-termination using the RADIUS protocol. The survival server performs authentication or query requests when authentication survivability is enabled, *and* one of the following is true:

1. All servers are out of service in the server group if fail-through is disabled.
2. All in-service servers failed the authentication and at least one server is out of service when fail-through is enabled.

All access credentials and key reply attributes saved in the local survival server remain in the system until they expire. The system-wide lifetime parameter **auth-survivability cache-lifetime** has a range from 1 to 168 hours, and a default value of 24 hours. Expired user credential attributes and key reply attributes stored in the survival server cache are purged every 10 minutes.



Best practices is to import a customer server certificate into the managed device and assign it to the local survival server.

The survival server can store the following types of client data:

- Client username
- Encrypted Passwords. For PAP authentication, the survival server receives the password provided by the client and then stores the encrypted SHA-1 hashed value of the password.
- EAP indicator: When employing 802.1X with disabled termination using EAP-TLS, the EAP indicator is stored.
- The CN lookup *EXIST* indicator

Supported Client and Authentication Types

The following combination of clients and authentication types are supported with the authentication survivability feature see the table below:

Table 45: *Clients and Supported Authentication Types*

Clients	Authentication Methods
Captive Portal clients	PAP
802.1X clients	■ *Termination disabled*: Extensible Authentication Protocol-Transport Layer Security with an external RADIUS server ■ *Termination enabled*: EAP-TLS with CN lookup with an external authentication server
External Captive Portal clients using the XML-API	PAP
MAC-based Authentication clients	PAP
VPN clients	■ PAP with an external authentication server ■ CN lookup with an external authentication server
VIA and other VPN clients	PAP method and CN lookup
Wireless Internet Service Provider roaming clients	PAP

Supported Key Reply Attributes

The following key reply attributes are supported:

- ARUBA_NAMED_VLAN
- ARUBA_NO_DHCP_FINGERPRINT
- ARUBA_ROLE
- ARUBA_VLAN
- MS_TUNNEL_MEDIUM_TYPE
- MS_TUNNEL_PRIVATE_GROUP_ID
- MS_TUNNEL_TYPE
- PW_SESSION_TIMEOUT
- PW_USER_NAME

Feature Restrictions and Limitations

The authentication survivability feature has the following support restrictions:

- The Survival Server cache database is station-based (thus, the MAC address is the key), so authentication survivability is not supported for any station with a zero MAC address.
- For a client using EAP-TLS, you must install the issuer certificate of the Survival Server certificate as a TrustedCA certificate in the client station.
- For an 802.1X client using EAP-TLS that does not terminate at the managed device, the issuer certificate for the client certificate must be imported as a TrustedCA or an intermediateCA certificate at the managed device—just as the same certificate must be installed at the terminating External RADIUS server.
- The Survival Server does not support the OCSP nor the CRL for EAP-TLS.
- Authentication survivability will not activate if Authentication Server Dead Time is configured as 0.

To configure Authentication Server Dead Time, on the managed device, navigate to: **Configuration** > **SECURITY** > **Authentication** > **Advanced** > **Authentication Timers** > **Authentication ServerDeadTime (min)**.

Captive Portal Authentication Workflow

This section describes the authentication procedures for Captive Portal clients, both when the branch authentication servers are available and when they are not available. When the authentication servers are not available, the Survival Server takes over the handling of authentication requests.

Captive Portal Client Authentication Using PAP

Table 46 describes what occurs for Captive Portal clients using PAP as the authentication method.

Table 46: *Captive Portal Authentication Using PAP*

When Authentication Servers Are Available	When Authentication Servers Are Not Available
▪ If authentication succeeds, the associated access credential with an encrypted SHA-1 hash of the password and Key Reply attributes are stored in the Survival Server database. ▪ If authentication fails, the associated access credential and Key Reply attributes associated with the PAP method (if they exist) are deleted from the Survival Server database.	When no in-service server in the associated server group is available, the Survival Server is used to authenticate the Captive portal client using PAP. The Survival Server uses the previously stored unexpired access credential to perform authentication and, upon successful authentication, returns the previously stored Key Reply attributes.

External Captive Portal Client Authentication Using the XML-API

Table 47 describes the authentication procedures for External Captive Portal clients using the XML-API, both when the branch authentication servers are available and when they are not available. When the authentication servers are not available, the Survival Server takes over the handling of authentication requests.

Table 47: *Captive Portal Authentication Using XML-API*

When Authentication Servers Are Available	When Authentication Servers Are Not Available
For authentication requests from an External Captive Portal using the XML-API, PAP is used to authenticate these requests with an external authentication server. ▪ If authentication succeeds, the associated access credential with an encrypted SHA-1 hash of the password and Key Reply attributes are stored in the Survival Server database. ▪ If authentication fails, the associated access credential and Key Reply attributes associated with the PAP method (if they exist) are deleted from the Survival Server database.	When no in-service server in the associated server group is available, the Survival Server is used to authenticate the Captive portal client using PAP. The Survival Server uses the previously stored unexpired access credential to perform authentication and, upon successful authentication, returns the previously stored Key Reply attributes.

802.1X Authentication Workflow

This section describes the authentication procedures for 802.1X clients with termination at an External RADIUS server, or at the controller.

Table 48: *802.1X Authentication Terminating at an External Server*

When Authentication Servers Are Available	When Authentication Servers Are Not Available
For an 802.1X client that terminates at an external RADIUS server using EAP-TLS: ■ If authentication is accepted, the associated access credential with the *EAP-TLS* indicator, in addition to the Key Reply attributes, are stored in the Survival Server database. ■ If authentication is rejected, the associated access credential and Key Reply attributes associated with the EAP-TLS method (if they exist) are deleted from the Survival Server database.	When there is no available in-service server in the associated server group, the Survival Server terminates and authenticates 802.1X clients using EAP-TLS. The Survival Server uses the previously stored unexpired access credential to perform authentication and, upon successful authentication, returns the previously stored Key Reply attributes. In this case, the client station must be configured to accept the server certificate assigned to the Survival Server.

For an 802.1X client for which termination is enabled at the managed device using EAP-TLS with CN lookup, a query request about the Common Name is sent to the external authentication server.



The external authentication server can be either a RADIUS server or an LDAP server.

Table 49: *802.1X Client Authentication Using EAP_TLS with CN Lookup*

When Authentication Servers Are Available	When Authentication Servers Are Not Available
■ If the query succeeds, the associated access credential with a returned indicator of *EXIST*, plus the Key Reply attributes, are stored in the Survival Server database. ■ If the query fails, the associated access credential and Key Reply attributes associated with the Query method (if they exist) are deleted from the Survival Server database.	When there is no available in-service server in the associated server group, the Survival Server performs CN lookup for 802.1X clients for which termination is enabled at the managed device using EAP-TLS. The Survival Server returns previously stored Key Reply attributes as long as the client with the *EXIST* indicator is in the Survival Server database.

MAC Authentication Workflow

This section describes the authentication procedures for clients.

Table 50: *MAC-Based Client Authentication Using PAP*

When Authentication Servers Are Available	When Authentication Servers Are Not Available
▪ If authentication succeeds, the associated access credential, along with an encrypted SHA-1 hash of the password and Key Reply attributes, are stored in the Survival Server database. ▪ If authentication fails, the associated access credential and Key Reply attributes associated with the PAP method (if they exist) are deleted from the Survival Server database.	When there is no available in-service server in the associated server group, the Survival Server authenticates the MAC-based authentication client using PAP. The Survival Server returns previously stored Key Reply attributes as long as the client with the *EXIST* indicator is in the Survival Server database.

WISPr Authentication

This section describes the authentication procedures for WISPr clients, both when the branch authentication servers are available and when they are not available. When the authentication servers are not available, the Survival Server takes over the handling of authentication requests.



The external authentication server can be either a RADIUS server or an LDAP server.

Table 51: *WISPr Authentication Using PAP*

When Authentication Servers Are Available	When Authentication Servers Are Not Available
For a WISPr client authenticated by an external server using PAP: ▪ If authentication succeeds, the associated access credential, along with an encrypted SHA-1 hash of the password and Key Reply attributes, are stored in the Survival Server database. ▪ If authentication fails, the associated access credential and Key Reply attributes (if they exist) associated with the PAP method are deleted from the Survival Server database.	When there is no available in-service server in the associated server group, the Survival Server authenticates the WISPr client using PAP. Upon successful authentication, the Survival Server uses the previously stored unexpired credential to perform authentication, and returns the previously stored Key Reply attributes .

Managed Device WAN Dashboard

The **WAN** dashboard, in the **Monitoring** section of the WebUI, is the default landing page for a managed device with uplinks defined via the uplink manager. Starting with ArubaOS 8.1.0.0, the WAN dashboard also appears in the Mobility Master WebUI when a branch office controller is selected in the network hierarchy.

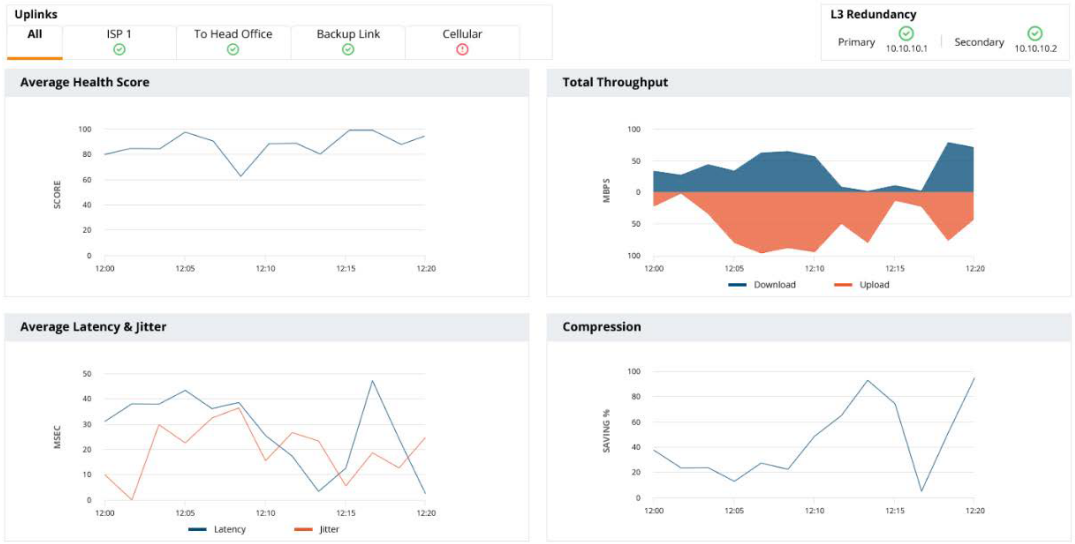


For more information on defining a WAN uplink, see Uplink Monitoring and Load Balancing on page 247 . For information on enabling and the uplink health check features, see Health Check Services for Managed Devices on page 244.

The WAN dashboard provides the WAN summary details for uplink VLANs, and contains the following tables:

- **Uplinks**: This section displays the **link** status and **WAN** status for VLANs monitored using the uplink manager utility. For each VLAN, the green check mark icon indicates an up status and red down arrow represents a down status for the link and WAN. The uplink health-check feature is disabled by default on managed devices. If it is enabled , the WAN status link will appear with a yellow icon, indicating that this feature is in an error state.
- **Health Score**: The health score rates the health of the uplink on a scale of 1-5, with score of 1 being lower quality and a score of 5 being the highest quality.
- **Throughput**: Displays the inbound and outbound traffic rates for the selected uplink.
- **Latency & Jitter**: Jitter is a variation in the delay times of received packets. If the **jitter measurement** option is enabled in the uplink manger, the uplink manger uses UDP packets on UDP port 4500 to measure jitter on the WAN links, and includes jitter statistics in the uplink quality calculations. Jitter statistics will not be measured if jitter measurement is not enabled in the uplink manager settings.
- **Aggregate Compression**: Displays the aggregate percentage compression on all VLANs with the compression feature enabled.

Figure 29 *WAN Monitoring Dashboard*



Using ZTP to Provision a Managed Device

When a factory-default controller boots, it starts the auto-provisioning process. The following sections describe the provisioning workflow, and the process to prepare your network for ZTP for a managed device.

When a managed device establishes an HTTPS connection to the Activate server and requests provisioning information, the Activate server authenticates the managed device and provides that device with provisioning information, including the IP address of its Mobility Master and secondary Mobility Master, and its country code.

If the managed device is unsuccessful in retrieving the provisioning parameters from Activate, it will retry in 30 seconds. The managed device will keep trying to retrieve the provisioning parameters from Activate until it is successful, or the administrator initiates Mini-Setup or Full-Setup provisioning.

Before you can use Activate to associate a managed device to Mobility Master, you must configure Activate with additional device settings for each managed device and Mobility Master, create a folder for those local devices, then assign a provisioning rule to the folder that associates the managed devices to

a specified master and configuration node. Use the following procedures to configure device details for the Mobility Master and managed devices, create folders, and define the provisioning rule.

Upgrading a Legacy Device via Activate

Starting with ArubaOS 8.1.0.0, a factory-default controller running ArubaOS 6.0.0.0 can use Activate Zero-Touch Provisioning to upgrade its software as part of the provisioning process. If Activate detects that a factory-default managed device running ArubaOS 6.x has been assigned a **Managed Device to Master Controller** provisioning rule, Activate will automatically send that managed device the information it needs to automatically download and upgrade to the latest version of ArubaOS.

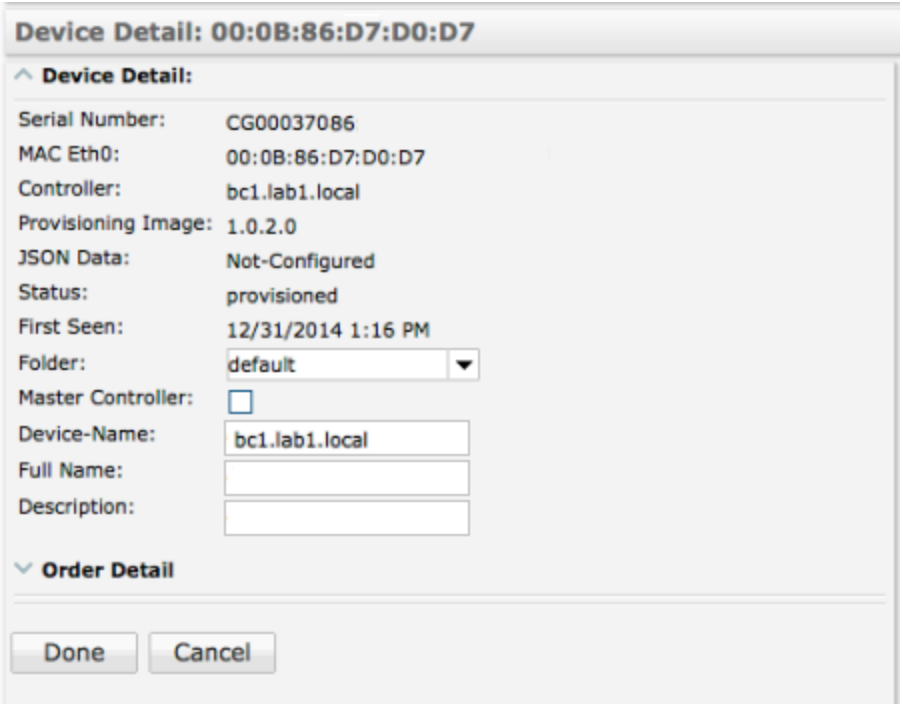
Configuring Device details for a Managed Device

When you place an order for a controller, that device appears in the Activate **Devices** list displaying the preconfigured settings for its serial number, MAC address, and software image. Before you can add a managed device to a whitelist, you must use the Activate interface to assign a name to each managed device, and use the Activate interface to identify the Mobility Master in a managed device deployment.

The following procedure describes how to configure managed device or Mobility Master device settings using Activate:

1. Click the **Devices** icon at the top of the page to display the **Devices** page.
2. Select a managed device or Mobility Master from the **Devices** list. If the list is very large, you can click the **filter** icon by any **Devices** list column heading and choose which entries to display, then select the managed device from the smaller, filtered list.
3. If the device will be used as the Mobility Master, select the **Master Controller** check box.
4. In the **Device Detail** section of the **Devices** page, enter the following values:
   - **Device name**: (Required) an IP address or fully-qualified domain name for the managed device or Mobility Master
   - **Full name**: (Optional) a user-friendly name for the device
   - **Description**: (Optional) a short text string describing the device
5. Click **Done** to save your settings.

Figure 30 *Device Details for a Managed Device*



Creating a New Managed Device Folder

Associate multiple managed devices to the same Mobility Master by moving those managed devices into a single Activate folder.

NOTE

A folder can contain only one model of managed device, using the same country code and mapping to the same configuration node. Different folders need to be created for managed devices of different model types, or that use a different country code or local configuration group.

The following procedure describes how to add a new folder to the **Folders** list:

1. Click the **Setup** icon to display the **Setup** page.
2. Click the **New** link in the title bar of the Folders list. The **Create a New Folder** window appears.
3. Enter the following information for the folder:
   - **Name**: Name of the new managed device folder. The folder name must be 100 characters or less, and cannot include the characters **?**, **#** or **&**.
   - **Parent**: The parent folder for the new folder. The new folder will be created under the selected parent.
   - **Notes**: (Optional) Use this field to add any additional notes about the folder.
4. Click **Done** to save the new folder.

Configuring the Provisioning Rule

A folder can only have one provisioning profile configured within it and the provisioning profile can only reference one configuration node. Consequently, it is necessary to create a folder and associate the provisioning rule for each group of managed devices that share a common configuration node.

The following procedure describes how to create a new provisioning rule for the new managed device folder:

1. Click the **Setup** icon to display the **Setup** page.
2. In the folders section of the **Setup** page, select the new managed device folder.
3. Click the **New** link in the title bar of the **Rules** list. The **Create a New Rule** window appears at the bottom of the page. Enter a value for each required field, then click **Done** to save your settings.

Figure 31 *New Provisioning Rule*

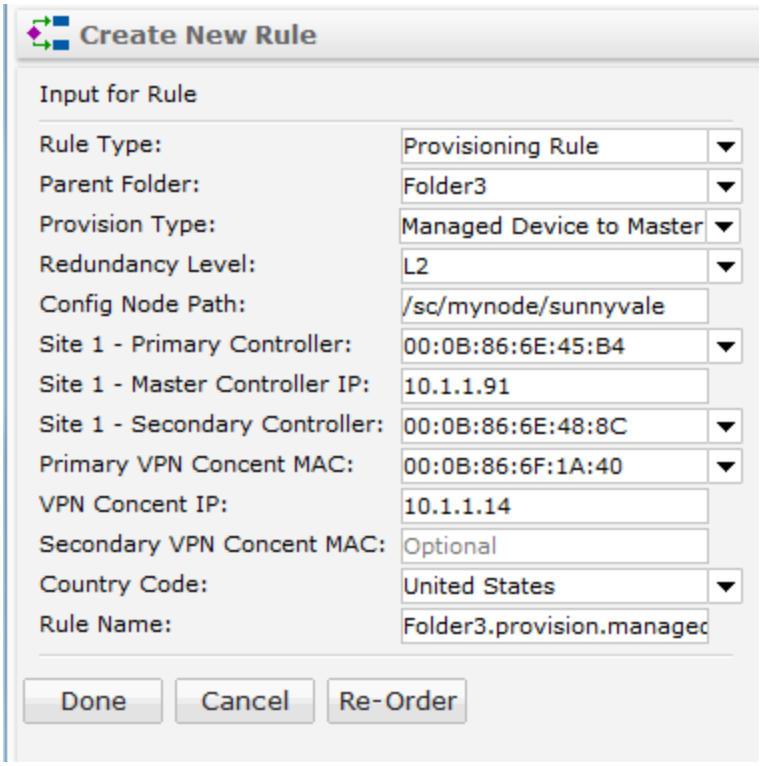


Table 52: *Provisioning Rule Configuration Settings*

Provisioning Rule Setting	Description
Rule Type	Click the **Rule Type** drop-down list, and select **Provisioning Rule**.
Parent Folder	Select the folder to which this provisioning rule applies.
Provision Type	Select the **Managed Device to Master Controller** rule type.
Redundancy Level	Select **No Redundancy** to configure just a single Mobility Master, choose **L2** redundancy to define a local backup at the same site as the Mobility Master or select **L3** to define an additional primary and backup Mobility Master at a different location than the main primary and backup Mobility Master pair. **NOTE:** If you select the L3 option, you must configure a Master Mobility Master and Secondary Mobility Master for Site 1 and Site 2.
Primary Controller	MAC address of the primary Mobility Master. Activate sends a managed device whitelist with information about the managed devices in this folder to the Mobility Master with this MAC address.

Provisioning Rule Setting	Description
Master Controller IP	Enter the IP address used to access Mobility Master or the primary/backup Mobility Master pair.
Secondary Controller	(Optional for Layer-2 or Layer-3 redundancy) MAC address of a backup Mobility Master, for deployments that require layer-2 or Layer-3 redundancy.
VPN Concentrator MAC	The MAC address of the managed device (or other device) that terminates VPN tunnels to the datacenter.
VPN Concentrator IP	The IP address of the managed device (or other device) that terminates VPN tunnels to the datacenter.
Country Code	Select a country code to be assigned to the managed devices in this folder.
Local Config Group	Enter the name of a local configuration group to assign that group of local configuration settings to the managed devices in this folder.

Moving a Managed Device to the New Folder

The following procedure describes how to assign one or more managed devices to a folder:

1. Click the **Devices** icon at the top of the page to display the **Devices** page.
2. Click the **filter** icon by any Devices list column heading and choose which entries to display. You can repeat this step and filter the list by multiple criteria types until the **Devices** list shows only those devices you want to move to a new folder.
3. Click the **Move to Folder** button at the top of the **Devices** page. A drop-down window appears, displaying with all folder names.
4. Select the destination folder for the devices.
5. A confirmation window appears, showing the total number of devices that will be moved.
6. Click **OK** to confirm the change, or click **Cancel** to cancel the move.

You can also assign an individual device to a new folder by selecting that device from the **Devices** list and manually changing its parent folder in the **Device Details** window.

Retrieval of a Managed Device Whitelist from Activate

Activate may be configured to supply the list of managed devices to the Mobility Master to be added to the whitelist.

The Mobility Master sends a query to Activate every hour. To initiate an immediate query to Activate, access the Mobility Master through CLI and issue the command “activate sync.”

When the Mobility Master sends the query to Activate, Activate searches for all provisioning rules of the type **managed node to master controller** that include the MAC address of this Mobility Master in the primary controller field.

Activate Interface Communication

The managed device and the Mobility Master interact with the Activate server to receive information about each other. Once the Activate server is properly configured with the appropriate folders and provisioning rules, Activate automatically manages the relationship between Mobility Master and all the managed devices associated with that master.

The Mobility Master regularly contacts the Activate server to get a list of its associated managed devices. Managed devices interact with the Activate server to learn about their role, Mobility Master information, and their regulatory domain. The Mobility Master sends its own information and not managed device information. Activate reuses information in the AP-information field for controller interactions between Mobility Master and managed devices.

The following steps describe how Mobility Master retrieves the whitelist database from the Activate server.

1. The Mobility Master sends an initial post with a keepalive connection type that includes the following information:
   - type = Provision update
   - mode = controller
   - a session ID
   - AP information that includes <serial number>, <mac-address>, <model>
2. Activate responds with the following information:
   - type = provision update
   - an Activate-assigned session ID
   - status
   - connection = keep alive.
3. The Mobility Master then sends a second POST with 'close' connection type with the following information:
   - type = provision update,
   - the session ID received from Activate,
   - Device information that includes <serial number>, <mac-address>, <model>
   - certificate length
   - signed certificate
   - device certificate
4. Activate then responds with the following information:
   - type = provision update,
   - the same session ID that Activate assigned in the first response
   - status = success or failure
   - mode = master
   - the list of managed devices from the whitelist database, where each list entry contains a <mac-address>,<serial number>,<model>,<mode>,<hostname>, and <config group>

Using ZTP with DHCP to Provision a Managed Device

The auto-provisioning process begins when a factory-default controller boots up. The following section describes the provisioning work flow, and also details the process to prepare your network for ZTP using DHCP.

In the absence of an Activate server, DHCP servers aid the managed devices to get information about the Mobility Master. The information required for provisioning managed devices is obtained from a DHCPv4 or DHCPv6 server.

Option 43 of DHCPv4 contains information about the Mobility Master to the managed devices. Similarly, for DHCPv6 Option 16 provides vendor related information and Option 17 provides information such as master IPv6 address, VPNC information and so on.

Following are the list of supported topologies:

- VMM with VPNC
- HMM with VPNC
- HMM without VPNC

VPNC must be a hardware controller and not a virtual machine.



In scenarios where both Activate and DHCP Option 43 are available, DHCP option 43 takes precedence over the Activate server. If an Activate server has to be used, then Option 43 should be removed from the DHCP server.

This feature also supports L2 Mobility Master Redundancy scenarios, where the managed device gets information about the primary Mobility Master and standby Mobility Master.

In VPNC scenarios, the managed devices get information related to primary Mobility Master, standbyMobility Master, Primary VPNC, and standby VPNC.

IPv4 Deployment Scenario

Option 43 of DHCPv4 contains the following information required to provision a managed device:

- masterip, country-code, master-mac1 (No L2 redundant Master)
- masterip, country-code, master-mac1, master-mac2 (L2 Redundant Master)
- masterip, country-code, vpnc ip, vpnc-mac1 (No L2 , Redundant VPNC)
- masterip, country-code, vpnc ip, vpnc-mac1, vpnc-mac2 (L2 Redundant VPNC)

Enter the details using one of the formats given below:

```
mip=10.9.186.001, mm1=aa:aa:aa:aa:aa:aa, cc=US
mip = 10.9.195.111 , cc= US,   vm2= 00:0C:20:C9:10:34 , vm1= 00:0C:29:B1:05:56A,
vip=10.45.12.111
```

Following is an example of a DHCPv4 configuration used for ISC DHCP server software:

```
subnet 10.3.91.0 netmask 255.255.255.0 {
option vendor-class-identifier  "ArubaMC";
option vendor-encapsulated-options "mip = 10.9.196.160 , cc= US, vm2= 00:0C:29:B9:20:64 ,
vm1= 00:0C:29:B9:20:5A, vip=10.45.34.187";
option domain-name-servers 10.1.10.10;
option routers 10.3.91.254;
range 10.3.91.2 10.3.91.253;
authoritative;
}
```

IPv6 Deployment Scenario

For DHCPv6, Option 16 contains Vendor Class Identifier (VCI), which is a text string that uniquely identifies a type of vendor device and Option 17 contains the following information required to provision a managed device:

- Master IPv4 (Optional)
- Master IPv6
- VPNC IPv6
- Primary Master MAC
- Redundant Master MAC
- Primary VPNC MAC
- Redundant VPNC MAC
- Country Code

Enter the details using the format given below:

```
mip=10.9.199.120,vip=10.9.199.100,vip6=2002:1:1:101::20,vm1=00:1a:1e:01:10:b0,mip6=2002:1
:1:101::361,cc=US,mm1=00:1a:1e:01:7d:c0
```

Following is an example of a DHCPv6 configuration used for ISC DHCP server software:

```
subnet6 2500:abcd:1234:dead::/64 {
default-lease-time 43200;
max-lease-time 43200;
option vendor-class-identifier  "ArubaMC";
option dhcp6.vendor-opts
“mip=10.9.199.120,vip=10.9.199.100,vip6=2002:1:1:101::20,vm1=00:1a:1e:01:10:b0,mip6=2002:
1:1:101::361,cc=US,mm1=00:1a:1e:01:7d:c0”
option dhcp6.name-servers  2111::1;
}
```

Health Check Services for Managed Devices

The health-check feature uses ping-probes to measure WAN availability and latency on selected uplinks. Based upon the results of this health-check information, the managed device can continue to use its primary uplink, or failover to a backup link. Latency is calculated based on the round-trip time of ping responses. You must define an uplink interface via the uplink manager and enable the health check feature before the results of this health check appear in the **WAN** section of the Monitoring Dashboard.



For more information on the WAN Dashboard, see WAN on page 426.

ArubaOS supports policy-based routing on each uplink interface, which allows you specify the next hop to which packets are routed. ArubaOS supports multiple next-hop lists, to ensure connectivity in the event that a device on the list becomes unreachable. If you are using Policy Based Routing, you can define global ping settings for all next-hop list destinations.

The **Health Check** section of the **Configuration > Services > WAN** tab allows you to configure probe measurement settings ping probe settings for the primary **WAN** uplink on the managed device, as well as for next hop links used by the policy-based routing feature

Table 53: *WAN Health Check Settings*

Parameter	Description
Health Check	Click this check box to enable the health check features.
Remote Host IP/FQDN	IP address or FQDN of a remote host to which the managed device is connected. The WAN health check feature will check the connectivity to the managed device uplink to this device.
WAN	
Probe Mode	Click the **Probe Mode** drop-down list and select **ping** or **UDP** to enable this feature.
Probe Interval (sec)	The **Probe Interval** field specifies the probe interval, in seconds. The WAN health-check feature sends the number of probes defined by the **Pocket Burst per Probe** parameter during each probe interval. To change the default interval of 10 seconds, enter a new value into this field.
Packet Burst Per Probe	The **Pocket Burst per Probe** field specifies the number of probes to be sent during the probe interval. To change the default value of 5 probes, enter a new value into this field.

Table 53: *WAN Health Check Settings*

Parameter	Description
Probe Retries	The number of times the managed device will attempt to resend a probe.
Jitter Measurement	If the health check feature is configured to use **UDP** probe mode, the WAN health-check feature can measure jitter on the connection to the remote host by sending and measuring packets at fixed intervals.
PBR	
Probe Mode	Click the **Probe Mode** drop-down list and select **ping** to enable this feature.
Probe Interval (sec)	The **Probe Interval** field specifies the probe interval, in seconds. The WAN health-check feature sends the number of probes defined by the **Pocket Burst per Probe** parameter during each probe interval. To change the default interval of 10 seconds, enter a new value into this field.
Packet Burst Per Probe	The **Pocket Burst per Probe** field specifies the number of probes to be sent during the probe interval. To change the default value of 5 probes, enter a new value into this field.
Probe Retries	The number of times the managed device will attempt to resend a probe. To change the default value of 3 retries, enter a new value into this field.

WAN Optimization Through IP Payload Compression

Data compression reduces the size of data frames that are transmitted over a network link, thereby reducing the time required to transmit the frame across the network. IP payload compression is one of the key features of the WAN bandwidth optimization solution, which is comprised of the following elements:

- IP Payload Compression
- Traffic Management and QoS



WAN optimization through IP payload compression is not supported in a 7205 controller.

The managed device can send traffic to destinations other than the corporate headquarters on the same link, so payload compression is enabled on the IPsec tunnel between the managed device and Mobility Master. Dynamic compression is used for the IP payload to achieve a high compression ratio. No compression is applied to data such as an embedded image file that might already be in a compressed format. Such data does not compress well, and may even increase in size.

The following procedure describes how to enable payload compression:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > WAN** tab.
2. Expand the **WAN Optimization** accordion.
3. Select the **Compression** option.
4. Click **Submit.**
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

WAN Interface Bandwidth Priorities

ArubaOS supports minimum bandwidth guarantees per traffic class, and allows critical delay-sensitive applications like voice and video to use more bandwidth and/or be scheduled with higher priority. Each interface can be associated with a scheduler profile that supports four queues with different priority levels. If you use session ACLs to define traffic policies on the managed device, you can use the scheduler profile to automatically associate these different priority levels assigned by these policies to a scheduler profile queue.



For information on creating a traffic policy that assigns 802.1p priority levels to a specific application or application type, see Firewall Policies on page 491

Each scheduler profile queue is assigned a priority level and one of the following scheduler discipline types:

- **Strict priority**: The queue service is based exclusively on the priority of the queue, where the lower priority queues are not serviced until the higher priority queue is clear. With this option, the highest level priority is guaranteed as much bandwidth as possible, but there can be phases where the 2nd, 3rd and 4th priority queues may receive little or no bandwidth.
- **Deficit Round Robin Weight**: The queue is assigned a percentage of available bandwidth.



You can define both strict priority and DDR Weight discipline types for a single scheduler profile.

The following procedure describes how to enable WLAN interface bandwidth priorities using the WAN scheduler feature:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Services > WAN** tab.
2. Expand the **WAN Scheduler** accordion.
3. Click + below the **WAN Scheduler Profiles** table to define a new scheduler profile.
   a. Enter a **Profile name**.
   b. In the **Priority** fields, enter one or more 802.1p priority levels (0-7) for each queue type. Each of the seven priority levels must be supported by one of the four queues.
   c. For each queue, click the **Scheduler Discipline** drop-down list and select the **Strict Priority** or **DDR Weight** discipline type. If you select the **DDR weight** option, enter the percentage of available bandwidth that should be made available to traffic in the selected queue. This field appears to the right of the **DDR weight** option.



If you configure both of strict priority and Deficit Round Robin weighted queues, the strict priority queues should be specified together continuously, followed by the Deficit Round Robin weighted queues. For example, if you want to specify two strict priority queues and two DDR weighted queues, configure queue 0 and 1 with the strict priority type, then configure queues 2 and 3 with a Deficit Round Robin priority type. You cannot alternate between strict priority and DDR weighted queues.

4. To assign the scheduler profile to a cellular or Gigabit Ethernet interface, click + below the **Assignments** table.
5. Click the **Ports** drop-down list to select an interface.
6. In the **Transmit Rate** field, enter the maximum transmit rate for the selected interface, in Mbps.

7. Create a firewall session policy that assigns a priority level to an application or application group. For details, see Firewall Policies on page 491

The following CLI commands enable WLAN interface bandwidth priorities:

```
(host)[node](config) #scheduler-profile <map-name> {priority-map <q0-q3> <que0-prio-list>} | {queue-weights <q0-q3> <percentage_weight>}
(host)[node](config) #interface cellular|gigabitethernet <slot/module/port> transmit max-rate rate mbits <mbps> scheduler-profile <profile>
(host)[node](config) #ip access-list session  any any app salesforce permit priority 3
```

Uplink Monitoring and Load Balancing

ArubaOS 8.5.0.0 and later versions do not support the uplink load balancing feature.

Wi-Fi Uplink

Starting from ArubaOS 8.5.0.0, Wi-Fi uplink is introduced to provide connectivity of AP to an external wireless network. The 3G/4G cellular uplink and the Wi-Fi uplink can be used to extend the connectivity to places where a wired uplink cannot be configured.

Wi-Fi uplink allows an AP running ArubaOS to connect to an external wireless network or a managed device by using a third-party AP, such as a Mi-Fi device or a smart phone running a hotspot. This requires the Aruba AP running ArubaOS to work as a standard Wi-Fi client. When the standard Wi-Fi client is used as an uplink, the AP requires MAC Address Translation (MAT) to bridge the traffic between wireless or wired users of the AP and the uplink network. Wi-Fi uplink can also be used to connect the AP to another Wi-Fi service, such as a hospital wireless network.



It is recommended to use Aruba mesh between one uplink Aruba AP and another Aruba AP. Wi-Fi uplink is used only when mesh is not suitable.

The ArubaOS AP must be provisioned with the necessary Wi-Fi uplink client parameters. After the AP reboots, it works as a standard client with the provisioned client parameters and connects with a Mi-Fi device or another AP to reach the managed device. The provisioned AP acts as both client and AP when it receives configurations from the managed device, which allows other wireless and wired clients to connect to the Aruba AP.



- Wi-Fi uplink is applicable to 802.11ax AP platforms and 802.11ac wave2 AP platforms except 340 Series access points.
- Wi-Fi uplink is not supported on 802.11ac wave1 AP platforms, including AP-204, AP-205, 210 Series, 220 Series, and 270 Series access points.
- Wi-Fi uplink is not supported on AP-555 in tri-radio mode.



The following sections describe how to configure a Wi-Fi Uplink profile and provision an AP with Wi-Fi uplink:

Configuring a Wi-Fi Uplink Profile

The following configuration conditions apply to Wi-Fi uplink:

- If the Wi-Fi uplink is used on 2.4 GHz or 5 GHz band, mesh or cellular uplink is disabled. The two links are mutually exclusive.
- To bind or unbind the Wi-Fi uplink on 2.4 GHz or 5 GHz band, reboot the AP.
- An AP provisioned with Wi-Fi uplink client parameters can failover to wired uplink and vice-versa, depending on the priority specified for Wi-Fi uplink and wired uplink. However, preemption is not allowed in this release.

The following procedure describes how to configure an AP with Wi-Fi uplink profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Expand the **AP** accordion.
3. Select **WiFi uplink**.
4. Select the Wi-Fi uplink profile that you want to edit or click **+** and enter a name into the **Profile Name** dialog box to create a new profile.



You can create upto 16 Wi-Fi uplink profiles with different priorities in an ap-group or ap-name.

1. Configure the Wi-Fi uplink profile settings described in Table 54.

Table 54: *Wi-Fi Uplink Profile Parameters*

Parameter	Description
General	
ESSID	Enter the required ESSID to which the client is associated.
BSSID	(Optional) Enter the required BSSID to which the client is associated.
Allowed band	Select one of the following radio band(s) on which the Wi-Fi uplink is used: ■ a: 802.11a band only (5 Ghz) ■ g: 802.11g band only (2.4 Ghz) ■ all: Both 802.11a and 802.11g bands (5 GHz and 2.4 GHz) Default: all
Security	
Encryption	Select one of the following data encryption types: ■ opensystem— No authentication or encryption. ■ static-wep— WEP with static keys. ■ personal— A wildcard mode that matches several PSK mode key management suites and cipher suites, including WPA-PSK-TKIP, WPA-PSK-AES, WPA2-PSK-TKIP and WPA2-PSK-AES. Default: opensystem.

Parameter	Description
WEP Key 1	Enter the first static WEP key associated with this key index. Can be 10 or 26 hex characters in length. Re-enter the key in the Retype text box.
WEP Key 2	Enter the second static WEP key associated with this key index. Can be 10 or 26 hex characters in length. Re-enter the key in the Retype text box.
WEP Key 3	Enter the third static WEP key associated with this key index. Can be 10 or 26 hex characters in length. Re-enter the key in the Retype text box.
WEP Key 4	Enter the fourth static WEP key associated with this key index. Can be 10 or 26 hex characters in length. Re-enter the key in the Retype text box.
WEP Transmit Key Index	Enter the key index to specify which static WEP key is to be used. Can be 1, 2, 3, or 4.
WPA Hexkey	Configure a WPA Pre-Shared Key (PSK). This key must be of 64 hexadecimal characters. Re-enter the key in the Retype text box.
WPA Passphrase	Configure the WPA password that generates the PSK. The passphrase must be between 8–63 characters, inclusive. Re-enter the password in the Retype text box.



When both WPA Hexkey and WPA Passphrase fields are configured, WPA Hexkey takes precedence.

The following CLI commands configure an AP with a Wi-Fi uplink profile:

```
(host)[mynode](config)# ap wifi-uplink-profile test-uplink
(host)[mynode](WiFi uplink profile "test-uplink")# essid uplink-new
(host)[mynode](WiFi uplink profile "test-uplink")# wpa-passphrase ********
(host)[mynode](WiFi uplink profile "test-uplink")# opmode personal
(host)[mynode](WiFi uplink profile "test-uplink")# exit
```

Provisioning an AP with Wi-Fi Uplink

The following procedure describes how to provision an Aruba AP with the Wi-Fi uplink:

1. In a **Managed Network** node hierarchy, navigate to **Configuration > Access Points**.
2. In the Campus APs tab, select the new AP from the **Campus APs** list, then click **Provision**.
3. In the AP provisioning section, click the **AP Group** drop-down list and select the AP group to which the Aruba AP should be assigned.
4. In **Controller discovery**, select **Use AP discovery protocol (ADP)** if you want to provide the AP with its managed device IP address, or select **Static** to manually define the managed device IP for that AP. If you select the **Static** option, you are prompted to enter the managed device's DNS name or IP address.
5. In **IP**, select **DHCP** if you have configured a DHCP server to provide the AP with the AP IP address, or select **Static** to manually define the AP IP address. If you select the **Static** option, you are prompted to enter the following information for the selected AP:
   - IPv4 address, netmask, internet gateway used by the AP, and DNS server.
   - IPv6 address, netmask, internet gateway used by the AP, and DNS server.

6. Select the **WiFi uplink** check box to enable Wi-Fi uplink on the AP.



When Wi-Fi uplink is enabled, at least one Wi-Fi uplink profile must be added to the AP group.

7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy Changes** to re-provision the AP.



You must re-provision the AP to enable Wi-Fi uplink profile on the AP. Re-provisioning the AP causes it to automatically reboot.

The following CLI commands configure the Wi-Fi uplink profile in the AP group:

```
(host)[mynode](config)# ap-group wfu-test
Warning: WiFi uplink profile will not take effect until an AP is reprovisioned
(host)[mynode](AP group "wfu-test")# wifi-uplink-profile test-uplink priority 1
```

The following CLI commands provision the AP with Wi-Fi uplink profile:

```
(host)[mynode](config)# provision-ap
(host)[mynode](config-submode)# read-bootinfo ip-addr 192.168.244.2
(host)[mynode](config-submode)# link-priority-wifi 10
(host)[mynode](config-submode)# ap-group wfu-test
(host)[mynode](config-submode)# wifi-uplink
(host)[mynode](config-submode)# reprovision ip-addr 192.168.244.2
```

Uplink Load Balancing

WAN traffic can be balanced across two or more active uplinks from a managed device to a VPN concentrator. The uplink load balancing feature supports both active and standby uplinks, so the traffic load is balanced across two wired uplinks, while the backup cellular uplink remains idle.

When a managed device has multiple active uplinks, uplink load balancing can modify the IKE parameters for the managed device to create multiple managed device/VPN concentrator IPsec tunnels, one on each uplink. Once multiple uplinks and IPsec tunnels are up, Layer-3 traffic can be load-balanced across these uplinks using specially created internal routing ACLs and nexthop lists.

Load Balancing ACLs

When uplink load balancing is enabled, any Layer-3 traffic session that is not associated to a manually defined routing ACL will be managed by two specially created, internal ACLs placed at the bottom of the routing ACL table; the editable ACL **uplink-lb-cfg-racl**, followed by the non-editable ACL **uplink-lb-sys-racl**.

Load Balancing Nexthop Lists

The uplink load balancing feature uses three special internally created nexthop lists:

Load-balance-gateways is used for load-balancing internet-bound traffic, and **load-balance-ipsecs** for managing encrypted traffic headed to the corporate headquarters. These nexthop lists include information about one nexthop gateway and one managed device / VPN concentrator IPsec tunnel for each uplink, which are added to these lists so all nexthops are considered active and are available for routing.

The third nexthop list created by this feature is **traditional-ipsecs**, which is created by the load balancing feature, and used by uplinks in active-standby mode to send control plane traffic from the managed device to the VPN concentrator.

Configuring the Uplink Manager

The following procedure describes how to disable or enable the uplink manager, and manage priorities for the wired and cellular connections. The uplink managed is enabled by default on managed device uplinks.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration> Services > WAN** tab
2. Expand the **Uplink** accordion.
3. Select the **Compression** check-box to enable WAN compression .
4. Select one of the following options from **Loadbalancing Mode** drop-down list.
   - **Hash based**: Hash-based load balancing uses information from the packets being sent (for example, the source IP address, destination IP address, protocol and port numbers to determine how to load balance that traffic).
   - **Round Robin**: Traffic is equally distributed to all the active uplinks.
   - **Session Count**: Traffic is balanced between the uplinks based upon the number of sessions managed by each link, so that the load for each active uplink stays within 5% of the other active uplinks.
5. Click **+** in the **Uplink VLANs** table and enter the following values to define a uplink VLAN for an uplink interface on the managed device.
   - **Link name**: Name of the uplink.
   - **Link type**: Type of the uplink. The options are MPLS, INET, LTE and Metro-Ethernet
   - **Interface VLAN ID**: VLAN ID number
   - **Operation State:** Select this check-box to disable or re-enable the VLAN.
   - **Use only as backup link**: Select this check-box to use this uplink only as a backup link.
   - **Weight**: If the uplink is using load balancing in **session** mode, this value defines the weight given to the uplink in an active/active uplink scenario. An uplink with a higher weight will be assigned more session traffic than an uplink with a lower weight. The supported range of values is 1-100.
6. Click **Submit.**
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following examples configure an uplink load-balancing solution using the Mobility Master command-line interface.

Step 1: Configure the VPN concentrator

If a managed device terminates a secure tunnel on a VPN concentrator, you can issue the **vpn-peer peer-mac** command on the VPN concentrator configuration to enable load balancing on secure uplinks between the VPN concentrator and a managed device.

The following CLI example enables uplinks between a managed device with the MAC address 01:00:5E:00:00:FF and a VPN concentrator, this automatically enables load balancing:

```
(host)[node](config) #vpn-peer peer-mac 01:00:5E:00:00:FF cert-auth factory-cert
```



.If the peer device is an x86 smanaged device. However, then configure the MAC address of the management interface of the However, if the peer device is a hardware platform, you must provide the MAC address of the VLAN interface of the managed device

Step 2: Configure Wired and Cellular Uplinks

The following CLI commands configure the wired and cellular uplinks in the uplink manager:

```
(host) [node] (config) #uplink cellular uplink-id
(host) [node] (config) #uplink wired vlan
```



The uplink manager and load-balancing settings are enabled automatically when you configure the cellular or wired uplinks.

Step 3: (Optional) Configure Load Balancing Settings

The following CLI commands configure additional uplink load-balancing settings.

```
(host) [node] (config) #uplink load-balance ?
mode                 Configure load-balancing mode
threshold-limits     Set threshold limits for load balancing
```

Hub and Spoke VPN Configuration

Mobility Master supports the hub and spoke VPN topology for Aruba branch office solutions. In this topology, one or more VPN routers (remote branches or spokes) communicate with a central VPN router (VPN Concentrator or hub) using a secured tunnel. The VPN Concentrator identifies the endpoints using the TPM certificates to establish the secured tunnel. This topology allows users at remote sites to access the main network and is best suited for networks where the traffic between the remote sites and the main network is predominant with minimal inter-site traffic.

Ensure to configure the VPN Concentrator and the managed devices at the branch locations to set up a hub and spoke VPN. You can configure 7200 Series Mobility Controllers as VPN Concentrators and 7000 Series Mobility Controllers as branch office devices.

This section includes the following topics:

- Whitelisting Managed Devices on VPN Concentrator on page 252
- Configuring VPN Tunnels on Managed Devices on page 253

Whitelisting Managed Devices on VPN Concentrator

In a hub and spoke VPN topology, where remote branches connect to the VPN Concentrator, newer branches are added in a staggered way. Each time a managed device is added to a branch, the branch information needs to be populated in the VPN Concentrator to whitelist the branch device. With large-scale deployments, this method is error prone and very cumbersome. The automatic whitelisting feature enables automating the process of whitelisting the branch devices to avoid extra configuration for each device at the headend.

For automatic whitelisting of managed devices in the VPN Concentrator, the authentication code method is used. In this method, the whitelisting of the device is achieved through the authentication token.

Configuring Passcode Based Whitelisting

You must configure the same VPN peer authentication passcode on the managed devices as well as the VPN Concentrator to whitelist the device in the database.

The following procedure describes how to whitelist a managed device automatically on a VPN Concentrator:

1. In the **Managed Network** node hierarchy, navigate to **Configuration> Services > VPN**.
2. Expand the **Hub and Spoke** accordion.
3. Enable the toggle switch **Hub and Spoke settings.**
4. Select **Hub (VPNC)** in **Deployment mode**.

5. Select **Automatic** in **Connection mode**.
6. Enter the same passphrase that is configured on the managed device in the **Passphrase** field for automatic whitelisting.
7. Select an encryption method from the **Encryption** drop-down list.
8. For **Custom Cert** encryption method, enter the **CA cert** and **Server cert** details.
9. Select **Route** or **Session** from the **ACL type** drop-down list based on your requirement and then select the appropriate ACL.
10. If you have overlapping uplink IP address across branches, then enter the branch pool details.
11. Click **Submit**.

The following CLI command configures the authenticate code on the Mobility Master which is used for automatic whitelisting of managed devices on a VPN concentrator where the same authenticate code is configured.

```
(host)[mynode] (config) #vpn-peer pass-code Aruba123 cert-auth factory-cert
```

Configuring MAC Address Based Whitelisting

The following procedure describes how to whitelist a managed device manually on a VPN Concentrator:

1. In the **Managed Network** node hierarchy, navigate to **Configuration> Services > VPN**.
2. Expand the **Hub and Spoke** accordion.
3. Enable the toggle switch **Hub and Spoke settings.**
4. Select **Hub (VPNC)** in **Deployment mode**.
5. Select **Manual** in **Connection mode**.
6. Click **+** from the **Branch Gateways** table to add the MAC address of the managed devices:
   - **MAC ADDRESS**—Enter the MAC address of the primary VPN Concentrator.
   - **ENCRYPTION**—Specify the encryption method. It can be **Factory Cert** or **Custom Cert**
   - **CA CERT**—Select the CA certificate for the custom certificate.
   - **SERVER CERT**—Select the server certificate for the custom certificate.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Configuring VPN Tunnels on Managed Devices

You can configure the managed devices to establish a VPN tunnel with the VPN Concentrator using one of the following methods:

- By configuring Auto-VPN to automatically establish a VPN tunnel with a VPN Concentrator by advertising the branch devices.
- By configuring a VPN endpoint for the managed devices to establish a VPN tunnel.

Configuring Auto-VPN on Managed Devices

The following procedure describes how to configure Auto-VPN using branch advertisement:

1. In the **Managed Network** node hierarchy, navigate to **Configuration> Services > VPN**.
2. Click **Hub and Spoke**.
3. Enable the toggle switch **Hub and Spoke settings.**
4. Select **Hub (VPNC)** in **Deployment mode**.

5. Select **Automatic** in **Connection mode**.
6. Enter the same passphrase that is configured on the VPN Concentrator for automatic whitelisting in the **Passphrase** field.
7. Re-enter the passphrase in **Confirm Passphrase** field.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Configuring VPN Endpoint for Managed Devices

The following procedure describes how to configure a specific VPN endpoint for the managed devices:

1. In the **Managed Network** node hierarchy, navigate to **Configuration> Services > VPN**.
2. Click **Hub and Spoke**.
3. Enable the toggle switch **Hub and Spoke settings.**
4. Select **Spoke (Branch Gateway)** in **Deployment mode**.
5. Select **Manual** in **Connection mode**.
6. Click **+** from the **Hubs** table to add the following VPN Concentrator hub information:
   - **Primary VPNC**—Enter the MAC address of the primary VPN Concentrator.
   - **Backup VPNC**—(Optional) Enter the MAC address of the backup VPN Concentrator.
   - **IP Address**—Enter the IP address of the VPN Concentrator.
   - **Source VLAN**—Specify the source VLAN of the managed device if more than one IP address is configured for the same VPN Concentrator.
   - **Encryption**—Specify the encryption method. It can be **Factory Cert** or **Custom Cert**.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

IP Routes Configuration

The managed devices and VPN Concentrator in a branch network must have IPv4 routes to determine how each device must reach Mobility Master and its VPN peers over any intermediate public or private IPv4 networks (underlay routes). Routes are also required to determine the internal networks that must be reached by the branch devices through the overlay VPN tunnels (overlay routes).

Underlay Routes

To reach WAN or the internet, the VPN Concentrators in data centers can use static routes. In case of private WAN deployments, the administrators can configure Open Shortest Path First (OSPF) routes.

managed devices, however, use the default routes obtained from service providers through DHCP or PPPoE. For private WAN deployments or MPLS routing, the administrators can configure static routes.

Overlay Routes

For overlay routes, the administrators can use IKEv2 extensions to dynamically learn networks from each connected branch. The routes can be populated in the forwarding table for each VPN Concentrator as static routes. These routes can also be redistributed into OSPF. The administrators can define static routes for each destination network and VPN Concentrator, and then configure VPN Concentrators to redistribute routes at different costs to prevent routing loops.

Configuring Static IP Routes

For overlay routing using static IP routes, ensure that you define static routes for each branch network and data center as follows:

- Static routes for each branch network must be defined on the router in the data center.
- Static routes for each branch network must be defined on the VPN Concentrator for each remote network, peer, and link.
- Static routes for each data center or a hub site must be defined for each managed device.

Creating a Static IP Route

To configure a static IP route, perform the following steps in the WebUI:

1. In the **Managed Network** node hierarchy, navigate to **Configuration> Interfaces > IP Routes** tab.
2. Expand **IP Routes** and click **+** to add a static route to a destination network or host.
3. Enter the IP address and netmask for the **Destination IP address** and **Destination network mask**, respectively.
4. Configure a forwarding setting:
   - **Using Forwarding Router Address**—Enter the next hop IP address in dotted decimal format (A.B.C.D). You can also enter the distance metric (cost) for this route. The cost prioritizes routing to the destination. The lower the cost, the higher the priority.
   - **Using IPsec Tunnel to VPNC**—Select the VPN Concentrator and the uplink to use. Select this option for a Hub and Spoke VPN configuration. For more information, see Hub and Spoke VPN Configuration on page 252.
   - **Using Site-to-Site IPsec**—Enter the IPsec map name to use in a static IPsec route map. Select this option for a site-to-site VPN. For more information, see Working with Site-to-Site VPNs on page 480.
   - **Using Null Interface**—Designate a null interface.
5. Specify a value for the **Cost**.
6. Click **Submit**.

Uplink Routing using Next-hop Lists

If the managed device uses policy-based routing to forward packets to a next-hop device, a next-hop list ensures that if the primary next-hop device becomes unreachable, the packets matching the policy can still reach their destination. ArubaOS now also allows IPv6 next-hop lists in policy-based routing. For more information on next-hop configuration, see Policy Based Routing on page 257.

Defining Next-hop Lists

The following procedure describes how to define a next-hop list:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration> Services> WAN** tab.
2. Expand the **Next Hop Configuration** accordion.
3. (Optional) In the **Health check probe interval** field, specify the probe interval, in seconds.
4. The default value is 10 seconds.
5. (Optional) In the **Pocket Burst per Probe** field, specify the number of probes to be sent during the probe interval.

6. The default value is 5 probes.
7. Click + below the **NextHop Lists** table to open the **NextHop** section that allows you to configure the following next-hop settings:

Table 55: *Managed Device Next-Hop Settings*

Parameter	Description
NextHop list name	Add a name for the new next-hop list. **NOTE:** You cannot use the same name for both IPv4 and IPv6 next-hop lists.
IP version	Select either **IPv4** or **IPv6** from the drop-down list, which you want to assign for the new next-hop list.
NextHops	IPv4 or IPv6 address of the next-hop device or the VLAN ID of the VLAN used by the next-hop device. If the VLAN gets an IPv4 address using DHCP, and the default gateway is determined by the VLAN interface, the gateway IP is used as the next-hop IP address. ▪ Click **+** to open the **Add IPv4 NextHop** pop-up window, if you selected **IPv4** option in the **IP version** field. In the **Add IPv4 NextHop** pop-up window, select one of the following radio buttons: • **IP**— Enter the IPv4 address and priority of the next-hop device In the **IP address** and **Priority** fields respectively. • **DHCP**— Enter the VLAN ID and priority of the next-hop device In the **VLAN ID** and **Priority** fields respectively. ▪ Click **+** to open the **Add IPv6 NextHop** pop-up window, if you selected **IPv6** option in the **IP version** field. In the **Add IPv6 NextHop** pop-up window, enter the IPv6 address and priority of the next-hop device in the **IPv6 address** and **Priority** fields. Use the optional **Priority** field to assign priority to next-hop device. The range is 1-255 and default value is 128. **NOTE:** You can configure a maximum of 16 next-hop devices for a next-hop list, and a maximum of 32 next-hop lists are currently supported. **NOTE:** You cannot configure IPv6 multicast, link-local, unspecified, loopback, and subnet anycast addresses as IPv6 next-hop addresses.
IPsec map name	A next-hop list may require policy-based redirection of traffic to different VPN tunnels. Select an IPsec map to redirect traffic through IPsec tunnels. Click **+** to open the **Add New IPsec Map** pop-up window. Select either **Using site-to-site IPSec** or **Using IPSec Tunnel to VPNC** option from the drop-down list of **Forward Settings** field, and specify the priority in the **Priority** field. **NOTE:** For IPv6 address, only **Using site-to-site IPSec** option is supported under **Forward Settings** field. If a managed device terminates a secure tunnel on a VPN concentrator, you can issue the **vpn-peer peer-mac** command on the VPN concentrator configuration to enable load balancing on secure uplinks between the VPN concentrator and a managed device. The following example enables uplinks between a managed device with the MAC address 01:00:5E:00:00:FF and a VPN concentrator, this automatically enables load balancing: `(host)[node](config) #vpn-peer peer-mac 01:00:5E:00:00:FF cert-auth factory-cert`

Parameter	Description
	NOTE: If the peer device is an x86 server, then configure the MAC address of the management interface of the managed device. However, if the peer device is a hardware platform, you must provide the MAC address of the VLAN interface of the managed device.
Preemptive-failover	If preemptive failover is disabled and the highest-priority device on the next-hop list is disabled, the new primary next-hop device remains the primary even when the original device comes back online.

8. Click **Submit.**
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box.
11. Click **Deploy Changes**.

Policy Based Routing

A policy-based routing rule is an ACL that can forward traffic as normal, or route traffic over a VPN tunnel specified by an IPsec map, routed to a next-hop router on a next-hop list, or redirected over an L3 GRE tunnel or tunnel group.

ArubaOS now also supports IPv6 address in policy-based routing rule.



A Policy Based Routing rule does not become active until it is applied to a VLAN interface or user role.

Associating PBR Rule with Managed Device

The following procedure describes how to associate a policy based routing rule with a managed device:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration> Services > WAN** tab.
2. Expand the **Policy-Based Routing** accordion.
3. Click **+** below the **Policies** table to create a new policy.
4. Enter the **Policy Name** in the **New Routing Policy** pop-up window and click **Submit**.
5. The policy type (route) is predefined in this window.
6. Select the policy created in the **Policies** table.
7. The **Policy > (policy name)** table is displayed.
8. Click **+** to add a new policy.
9. The **New Rule** pop-up window opens.
10. Select one of the following rule types:
    - **Access Control:** Applies the rule to all traffic, or traffic using a specific service, protocol, or TCP/UDP port or range of ports.
    - **Application**: Applies a rule to a traffic for an application or application category.



The **Application** rule type is not supported for IPv6 traffic.

11. Configure the rule parameters.

Table 56: *Policy Based Routing ACL Rule Parameters*

Field	Description
IP version	Select either **IPv4** or **IPv6** from the drop-down list to specify whether the policy applies to IPv4 or IPv6 traffic.
Source (required)	Source of the traffic, which can be one of the following: ■ **Any**: Acts as a wildcard and applies to any source address. ■ **User**: This refers to traffic from the wireless client. ■ **Host**: This refers to traffic from a specific host. When this option is chosen, you must configure the IP address of the host. ■ **Network**: This refers to a traffic that has a source IP from a subnet of IP addresses. When this option is chosen, you must configure the IP address and network mask of the subnet. ■ **Alias**: This refers to using an alias for a host or network. You configure the alias by navigating to the **Configuration > Advanced Services > Stateful Firewall > Destination** page. **NOTE:** When you select **IPv6** option in the **IP version** field, only **Any**, **Host**, and **Network** options are available as source of the traffic. **NOTE:** You cannot configure IPv6 multicast, link-local, unspecified, loopback, and subnet anycast addresses as IPv6 source addresses.
Destination (required)	Destination of the traffic, which can be configured in the same manner as source. **NOTE:** When you select **IPv6** option in the **IP version** field, only **Any**, **Host**, and **Network** options are available as destination of the traffic. **NOTE:** You cannot configure IPv6 multicast, link-local, unspecified, loopback, and subnet anycast addresses as IPv6 destination addresses.
IPv6 address	(Optional) Enter the IPv6 address to associate the policy to IPv6 traffic. **NOTE:** This field is visible only when you select **Host** under **Source** or **Destination** fields.
IPv6 netmask	(Optional) Enter the subnet mask for the IPv6 address. **NOTE:** This field is visible only when you select **Network** under **Source** or **Destination** fields.
Service/APP	If you are creating an **access control** rule, select a type of traffic, which can be one of the following: ■ protocol: Using this option, you specify a different layer 4 protocol (other than TCP/UDP) by configuring the IP protocol value. ■ any: This option specifies that this rule applies to any type of traffic. ■ service: Using this option, you use one of the pre-defined services (common protocols such as HTTPS, HTTP, and others) as the protocol to match for the rule to be applied. You can also specify a network service that you have manually configured. For details, see Creating a Network Service Alias on page 497. ■ tcp: A range of TCP port(s) that must be used by the traffic in order for the rule to be applied. ■ udp: A range of UDP port(s) hat must be used by the traffic in order for the rule to be applied.

Field	Description
	NOTE: : When you select **IPv6** option in the **IP version** field, only **Any** option is available as Service/App of the traffic.
Scope	If you are creating an **application** rule, select a type of traffic, which can be one of the following: ■ application: Create a rule that applies to a specific application type. Click the **Application** drop-down list and select an application type. ■ application category: Create a rule that applies to a specific application category. Click the **Application Category** drop-down list and select a category type.
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: ■ **Forward Regularly**: Packets are forwarded to their next destination without any changes. ■ **Forward to ipsec-map**: Packets are forwarded through an IPsec tunnel defined by the specified IPsec map. You must specify the position of the forwarding or routing rule. (1 is first, default is last) ■ **Forward to next-hop-list**: packets are forwarded to the highest priority active device on the selected next hop list. You must also specify the position of the forwarding or routing rule (1 is first, default is last). For more information on next-hop lists, see Uplink Routing using Next-hop Lists on page 255 ■ **Forward to tunnel**: Packets are forwarded through the tunnel with the specified tunnel ID. You must also specify the position of the forwarding or routing rule (1 is first, default is last). For more information on GRE tunnels, see GRE Tunnels on page 100. ■ **Forward to tunnel group**: Packets are forwarded through the active tunnel in a GRE tunnel group. You must also specify the position of the forwarding or routing rule (1 is first, default is last). For more information on tunnel groups, see GRE Tunnel Groups on page 107. **NOTE:** When you select **IPv6** option in the **IP version** field, only **Forward Regularly**, and **Route to next-hop-list** options are available.
Position	(Optional) Define a position for the rule in the ACL. Rules are processed according to their position numbers, and new rules are added at the end of an ACL by default. A position of 1 puts the rule at the top of the list. **NOTE:** The position that you select for an ACL rule is relative to either IPv4 or IPv6 policies.

12. Click **Submit**.
13. Click **Pending Changes**.
14. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Address Pool Management

Each managed device supports one or more client DHCP pools; a pool of IP addresses that can be assigned to clients associated to that managed device, or to the node itself. In addition to the DHCP pool, the Mobility Master also allows you to create separate pools of addresses a managed device can use to dynamically assign to its uplink VLANs, use for NAT translation, or use to create a GRE tunnel to the Mobility Master. These address pools are pushed out to each managed node when it comes up on the network. If a managed node is removed from the master, the IP addresses allocated to that managed device can be reused and reassigned to a new managed node.

ArubaOS supports the following pool types:

- DHCP Address Pools on page 260—When you create DHCP pool for a configuration group, that pool defines a set of IP addresses that can be assigned to client associated to managed devices in that group.
- VLAN Pools on page 261—Mobility Master must have a separate VLAN pool defined for each VLAN used by its managed device. A VLAN pool allocates a static, continuous block of multiple IP addresses to each managed device. The managed device acts as a DNS proxy server and dynamically assign IP addresses from its allocated pool to each AP or client on the VLAN.
- Tunnel Pools on page 262—The tunnel pool on a managed node defines a range of IP addresses that the managed node uses to create a GRE tunnel within the IPsec tunnel back to the Mobility Master. Unlike VLAN pools, which allocates multiple addresses to each managed node VLAN, the tunnel DHCP pool assigns a single tunnel IP address to each managed node.
- NAT Pools on page 262—Used by the managed device for source NAT translation. You can use a NAT pool to create a firewall policy rule to perform NAT on packets matching the rule.
- VPN Pools on page 263—The VPN pool defines a group of IP addresses assigned to VPN clients.

DHCP Address Pools

Use the **Configuration > Services > DHCP Server** page to configure a pool of DHCP addresses. The managed device can use one of the addresses from this pool for its own IP address, and/or assign addresses in the pool to clients associating to that node.

Configuring DHCP Address pool

The following procedure describes how to configure a DHCP address pool:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Services > DHCP > DHCP Server.**
2. Click **+** below the **Pool Configuration** table.
3. Define the following values for the pool, then click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Table 57: *DHCP Pool Configuration Parameters*

Parameter	Description
IP version	Assign IPv4 or IPv6 addresses
Pool Name	Give a name to the new address pool
Default router	IP address of the default router for the DHCP client. The client should be on the same subnetwork as the default router. You can specify up to eight IP addresses.
DNS Server	IP address of the DNS server. You can specify up to eight IP addresses. Multiple IP addresses must be separated by spaces.
Import from DHCP/PPPoE	Select this option to use the DNS server address obtained through PPPoE or DHCP.
Domain Name	Domain name to which the client belongs.
WINS	IP address of a NetBIOS Windows Internet Naming Service server. You can specify up to eight IP addresses. Multiple IP addresses must be separated by spaces.

Parameter	Description
Import from DHCP/PPPoE	Use the NetBIOS name server address obtained through PPPoE or DHCP.
Lease Days	The number of days that the assigned IP address is valid for the client.
Lease Hours	The number of hours that the assigned IP address is valid for the client.
Lease Minutes	The number of minutes that the assigned IP address is valid for the client.
Network IP Address Type	Choose **Static** to add a static IP address and netmask to the pool, or select **Dynamic** to define a range of addresses that the DHCP server may assign to clients. ■ If you select **Static**, enter an IP address and netmask. ■ If you select **Dynamic**, enter the starting and ending IP address for the address range, as well as the maximum number of hosts to be supported by the pool.
Option	Click **+** in **Option** to apply a client-specific option code and IP address or text string. See RFC 2132, "DHCP Options and BOOTP Vendor Extensions".

Assigning DHCP Address Pool to VLAN

The following procedure describes how to assign a DHCP address pool to a VLAN:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > interfaces > VLANs** tab.
2. In the **VLANs** table, select the name of the VLAN to which you want to assign the DHCP pool. A **VLANS > (selected VLAN)** table appears
3. Select the VLAN ID of the VLAN to use the address pool. The **Port Members** table opens.
4. In the **Port Members** table, select the **IPv4** subtab and expand the **IP Address Assignment** accordion.
5. For **IP assignment**, select **DHCP Pool**.
6. Click the **DHCP pool** drop-down list and select a DHCP to associate to the VLAN.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

VLAN Pools

You can create address pools for VLAN and assign them to the required VLAN interfaces. This topic includes the following sections:

- Creating Address Pools for VLANs on page 261
- Assigning Address Pool to VLAN on page 262

Creating Address Pools for VLANs

The following procedure describes how to create a VLAN pool for uplink interfaces on a managed device:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Interfaces> Pool Management** tab.
2. Expand the **VLAN Pools** accordion.
3. Click + below the **VLAN Pools** table to create a new VLAN pool

4. In the **Pool name** field, enter a name to the new pool.
5. In the **Start IP address** field, enter the IP address at the start of the range of addresses.
6. In the **End IP address** field, enter the IP address at the end of the range of addresses.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Assigning Address Pool to VLAN

The following procedure describes how to assign a VLAN address pool to a VLAN:

1. Navigate to **Configuration > interfaces > VLAN** tab.
2. In the **VLANs** table, select the name of the VLAN to which you want to assign the DHCP pool. A **VLANS > (selected VLAN)** table appears
3. Select the VLAN ID of the VLAN to use the address pool. The **Port Members** table opens.
4. In the **Port Members** table, select the IPv4 subtab.
5. For **IP assignment**, and select VLAN Pool.
6. Click the **VLAN Pool** drop-down list and select a DHCP to associate to the VLAN.

Tunnel Pools

The following procedure describes how to use tunnel pools to create a pool of IP addresses used by the managed device to create a GRE tunnel to the Mobility Master. Each managed device uses a single IP address from this pool.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces> Pool Management** tab.
2. Expand the **Tunnel Pools** accordion.
3. Click + below the **Tunnel Pools** table to create a new VLAN pool
4. In the **Pool name** field, enter a name to the new pool.
5. In the **Start IP address** field, enter the IP address at the start of the range of addresses.
6. In the **End IP address** field, enter the IP address at the end of the range of addresses.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following procedure describes how to associate a tunnel pool to a GRE tunnel:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Interfaces > GRE Tunnels** tab.
2. Select an entry in the **GRE Tunnels** table to associate a tunnel pool to that GRE tunnel.
3. In the **IPv4 Address Type** field, select the **Dynamic** option.
4. Click the **Dynamic IP Address Pool** drop-down list and select a tunnel pool.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

NAT Pools

The following procedure describes how to create a pool of addresses the managed device can use for Network Address Translation:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces> Pool Management** tab.
2. Expand the **VLAN Pools** section
3. Click **+** below the **VLAN Pools** table to create a new VLAN pool.
4. In the **Pool name** field, enter a name to the new pool.
5. In the **Start IP address** field, enter the IP address at the start of the range of addresses.
6. In the **End IP address** field, enter the IP address at the end of the range of addresses.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Nat pools can associated to firewall policy rules and VPN configurations.

- For information on creating a firewall policy rule that uses the NAT pool to performs NAT translation on matching packets, see Firewall Policies on page 491.
- To apply network address translation to VPN clients , navigate to **Configuration > Services > VPN > General VPN**, enable the **Source-NAT** option, then click the **NAT** drop-down list and select the NAT pool you just created.

VPN Pools

The following procedure describes how to create a pool of addresses used by VPN clients:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Services > VPN**.
2. Expand the **General VPN** accordion.
3. Click **+** below the **Address Pools** table to create a new VPN address pool
4. In the **Pool name** field, enter a name to the new pool.
5. In the **Start address IPV4 or V6** field, enter the IP address at the start of the range of addresses.
6. In the **End address IPV4 or V6** field, enter the IP address at the end of the range of addresses.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Configuring WAN Authentication Survivability

Enable WAN survivability for managed devices on your network by navigating to the **Configuration > Authentication> Advanced** tab, then selecting the Survivability tab.

The survivability settings on this tab are described in Table 58.



For additional information on WAN Authentication Survivability, including authentication workflows and supported client and authentication types see the WAN Authentication Survivability Overview on page 232.

Table 58: *WAN Authentication Survivability for a Managed Device*

Parameter	Description
Enable Auth-Survivability	This parameter controls whether to use the Survival Server when no other authentication servers in the server group are in-service. This parameter also controls whether to store the user access credential in the Survival Server when it is authenticated by an external RADIUS or LDAP server in the server group. Authentication Survivability is enabled or disabled at each managed device. This parameter is disabled by default. **NOTE:** Authentication Survivability will not activate if Authentication Server Dead Time is configured as 0. For more information on configuring Authentication Server Dead Time, see Configuring Authentication Timers on page 219.
Authentication Server Certificate	This parameter allows you to view the name of the server certificate used by the local Survival Server. The local Survival Server is provided with a default server certificate from ArubaOS. The customer server certificate must be imported into the managed device first, and then you can assign the server certificate to the local Survival Server.
Cache Lifetime (hrs)	This parameter specifies the lifetime in hours for the cached access credential in the local Survival Server. When the specified cache-lifetime expires, the cached access credential is deleted from the managed device. Configured authentication servers are put into the out-of-service state when authentication requests time out. The managed device picks the next server from the server group when the previous server times out or fails. When there are no more servers available from the server group, the local Survival Server processes the authentication request. When the client is authenticated with the local Survival Server, the previously stored Key Reply attributes are included in the RADIUS response. The Cache Lifetime range is from 1 to 168 hours. The default is 24 hours.
Certificate Type	Select the certificate to be used for client authentication.

Preventing WAN Link Failure on Virtual APs

In managed device deployments, the managed devices are connected across the WAN link from the Mobility Master to the RADIUS server. A WAN link outage will result in service outage as new users cannot be authenticated to 802.1X Virtual APs. This feature provides limited connectivity to managed devices even when the WAN link is down. To provide connectivity when the WAN link is down, open and PSK SSID Virtual APs are available at all times and the user can connect to these Virtual APs instead of the main 802.1X Virtual AP.



Currently, this feature is targeted for Campus APs in managed device deployments.

When all the WAN links are down, an AP management module in the controller updates the link state using the notification it receives from the health check manager. Depending on the link state, the new set of Virtual APs are made available to the users, ensuring minimum service depending on the deployment. The Virtual APs for WAN link failure feature can be configured using the Mobility Master WebUI or command-line interface.

The following procedure describes how to prevent the WAN link failure on virtual APs:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration> System> Profiles** tab.
2. In the **All Profiles** pane, expand the **Wireless LAN** menu.
3. Expand the **Virtual AP** menu.

4. Select an existing virtual AP profile.
5. Expand the **Advanced** accordion.
6. The **WAN Operation Mode** drop-down list supports the **primary**, **always**, and **backup** WAN modes. To enable WAN link failure, set this mode to **backup**.
7. Click **Submit.**
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Managed Device Integration with a Palo Alto Networks Portal

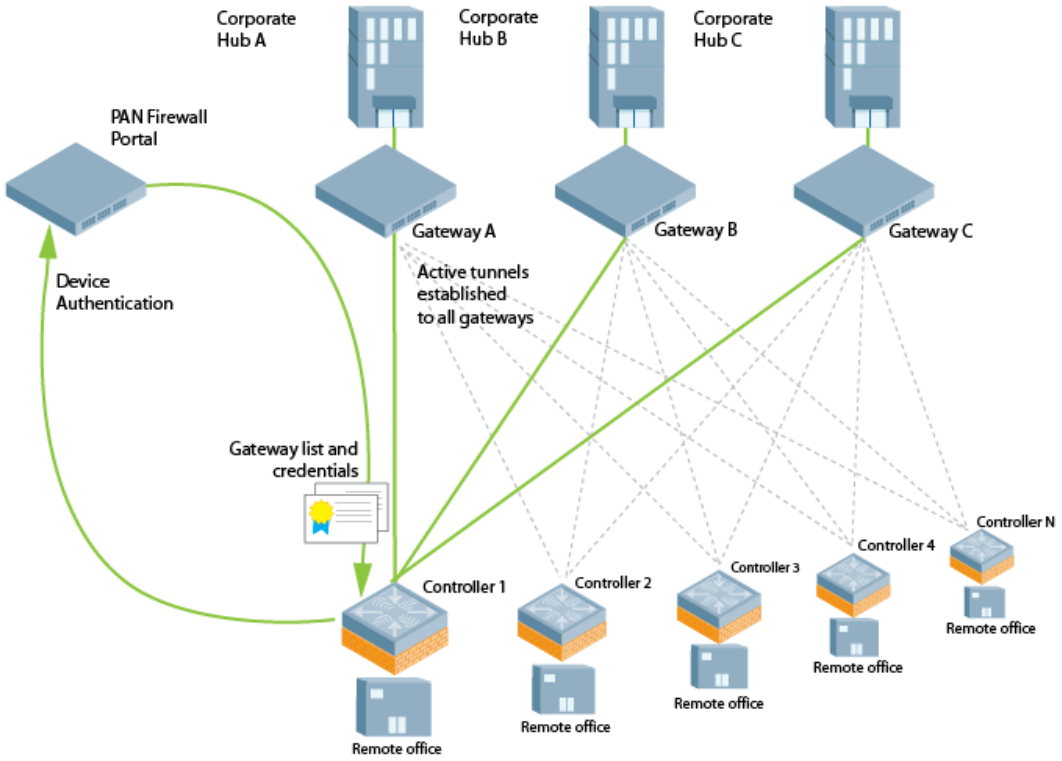
Managed devices can leverage their networks' existing Palo Alto infrastructure to access more advanced security services, including antivirus services, malware detection and seamless integration with the Palo Alto Networks WildFire™ cloud-based threat detection.

Overview

Enable Palo Alto firewall integration on Mobility Master to securely redirect internet inbound traffic from managed devices into the PAN firewall. Although this configuration setting can be used on a stand-alone Mobility Master, this feature can only be used in this types of deployments when used in conjunction with the Uplink VLAN manager feature.

The uplink VLAN manager is enabled by default on managed device uplinks. Stand-alone Mobility Masters using the PAN portal feature must enable the uplink VLAN manager using the **uplink** command in the Mobility Master command-line interface.

Figure 32 *Managed Device and PAN Firewall Integration*

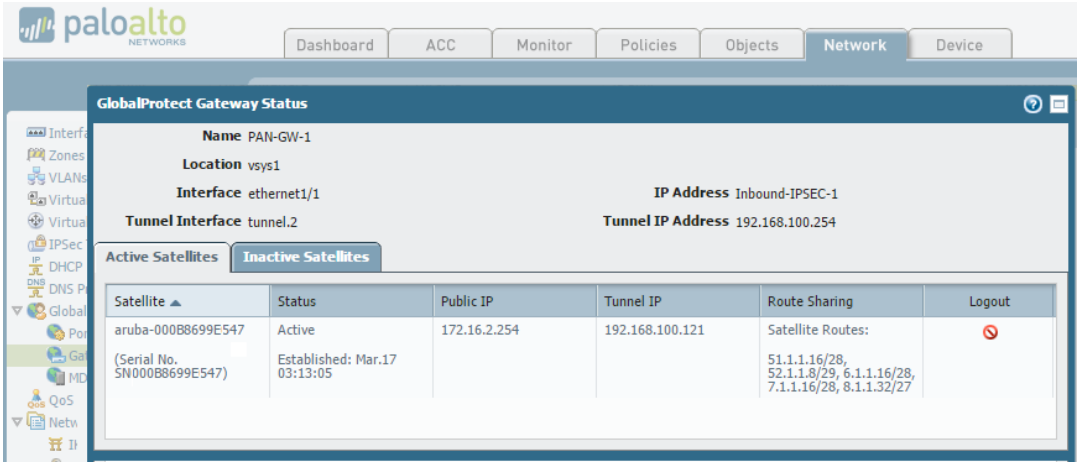


Integration Workflow

The following steps describe the work flow to integrate a managed device with a Palo Alto Networks Large-Scale VPN firewall.

1. Palo Alto Portal certificates are installed on Mobility Master, and the managed device is configured with the Palo Alto portal IP address or FQDN, Palo Alto certificate, and the username and password for device authentication using the **Configuration > Services > External Services > PAN Portal** section of the Mobility Master WebUI.
2. The managed device is provisioned via Aruba Activate and downloads its configuration (including Palo Alto Networks integration settings).
3. The Palo Alto portal may be configured with the device number (a text string comprised of the device serial number followed by its MAC address) of the managed device at each remote office site. This allows the managed device to bypass the username and password challenge to authenticate to the portal.
4. The managed device initiates a secure connection to the Palo Alto portal. Once the managed device is authenticated, the Palo Alto portal sends the managed device a list of PAN gateways and priority levels. Once the managed device is authenticated, that device appears in the PAN satellite list, as shown in the figure below.

Figure 33 *Palo Alto Networks Active Satellites List*



1. The managed device uses the Palo Alto Networks gateway list and credentials from the portal to contact all PAN gateways. Each PAN gateway sends the managed device information that allows the managed device to automatically create a secure IPsec tunnel and exchange branch subnet routes with each PAN gateway.
2. The managed device maintains a priority list of IPsec tunnels to each PAN gateway to enable failover in the event a PAN gateway becomes unreachable.
3. Policy-based routing ACL on the managed device selectively routes traffic to the PAN gateways.
4. Traffic redirected from the managed device is inspected via the Palo Alto Networks firewall.

Configuration Prerequisites

The Palo Alto Networks Large-Scale VPN framework can integrate with a managed device by establishing an IPsec tunnel between the firewall and the managed device. Integrating a Palo Alto Networks firewall with a managed device requires that all user traffic is routed, so it can be managed by a policy-based routing access control list.

The following certificate requirements must be fulfilled before the managed device can integrate with the Palo Alto Networks Large-Scale VPN framework:

- The Large-Scale VPN framework must be installed and active on your network. For more information on configuring Palo Alto Networks products, refer to the Palo Alto Networks Technical Documentation portal.
- The CA certificate used by the Palo Alto portal must be installed on Mobility Master, so that it can be pushed down to the managed device.
- On the PAN gateway devices, you must enable the **accept published routes** option, and the devices must install the server certificates derived from the management portal root CA.

In deployments with multiple PAN firewalls, you must configure the PAN management portal with a list of gateways and the priorities for each PAN gateway. Even if the PAN management portal uses serial number registration with preregistered serial numbers or MAC addresses, best practice is to configure LDAP, Radius, Kerberos or Local Database authentication as well. This allows a managed device to authenticate to the portal even if the portal does not recognize the managed device's MAC address.

Configuring PAN Portal settings

Pan portal settings must be defined via a managed device (/md) configuration. The Mobility Master configuration node (/mm) does not support PAN portal settings.

The following procedure describes how to configure PAN Portal settings:

1. From a **Managed Network** node hierarchy, navigate to the **Configuration > Services > External Services** tab.
2. Expand the **PAN Portal** accordion.
3. Define values for the configuration settings described in Table 59.

Table 59: *PAN Portal Settings*

Parameter	Description
Portal IP/FQDN	The IP address or FQDN of the portal.
Trusted Certificate	Specify the name of the self-signed or external CA certificate to establish an SSL connection to the portal.
User Name	Username to authenticate to the Palo Alto Networks portal.
Password	Password to authenticate to the Palo Alto Networks portal.

Chapter 14
802.1X Authentication

802.1X is an IEEE standard that provides an authentication framework for WLANs. 802.1X uses the 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-TLS, PEAP, and EAP-TTLS. These protocols allow the network to authenticate the client while also allowing the client to authenticate the network.

This chapter describes the following topics:

- Understanding 802.1X Authentication on page 268
- Configuring 802.1X Authentication on page 271
- Example Configurations on page 281
- Performing Advanced Configuration Options for 802.1X on page 299

Other types of authentication not discussed in this section can be found in the following sections of this guide:

- Captive portal authentication: Configuring Captive Portal Authentication Profiles on page 344
- VPN authentication: Planning a VPN Configuration on page 459
- MAC authentication: Configuring MAC-Based Authentication on page 222
- Stateful 802.1X, stateful NTLM, and WISPr authentication: Stateful and WISPr Authentication on page 311

Understanding 802.1X Authentication

802.1X authentication consists of three components:

- The *supplicant*, or client, is the device attempting to gain access to the network. You can configure the Aruba user-centric network to support 802.1X authentication for wired users and wireless users.
- The *authenticator* is the gatekeeper to the network and permits or denies access to the supplicants.
- The *Aruba managed device* 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 managed device.
- 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 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 IAS in Windows (see http://technet.microsoft.com/en-us/library/cc759077(WS.10).aspx).
- In Aruba user-centric networks, you can terminate the 802.1X authentication on the managed device. The managed device 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.

Starting from ArubaOS 8.4.0.0, the 802.1X authentication process is not part of the authentication manager. This enhancement allows the server to run multiple instances of new process for better performance.

This enhancement provides the following support:

- The logs that are part of 802.1X authentication are now listed in the 802.1X process instead of the authentication manager.
- When you enable the logs for the authentication server, the logs for the 802.1X process is automatically updated.

Starting from ArubaOS 8.4.0.0, the managed devices support EAP-TLS fragmentation as part of 802.1X authentication in non-termination mode. EAP-TLS fragmentation reduces RADIUS timeouts when:

- The size of an EAP packet exceeds 1500 bytes.
- A firewall exists between a managed device and an external authentication server but the external authentication server does not support RadSec.
- A firewall drops out-of-order IP fragments.
- A network uses active-active firewall and IP reassembly is incomplete

When enabled, configure a value of the IP MTU, with a minimal value of 576 bytes, to support EAP-TLS fragmentation.

Supported EAP Types

Following is the list of supported EAP types:

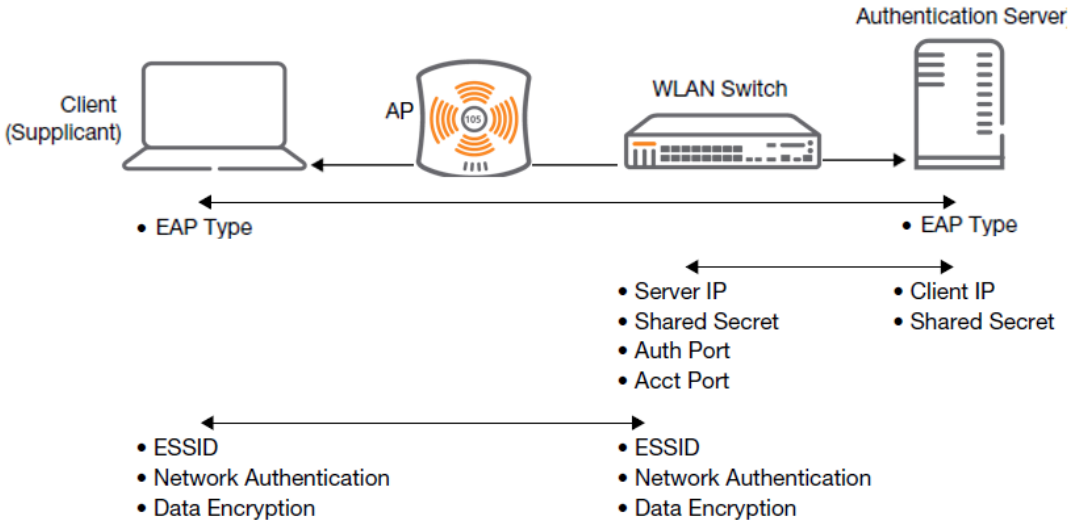
- PEAP — PEAP is an 802.1X authentication method that uses server-side public key certificates to authenticate clients with the server. The PEAP authentication creates an encrypted SSL or TLS tunnel between the client and the authentication server. The exchange of information is encrypted and stored in the tunnel to ensure that the user credentials are kept secure.
- EAP-GTC—The EAP-GTC type uses clear text method to exchange authentication controls between the client and the server. Since the authentication mechanism uses the one-time tokens (generated by the card), this method of credential exchange is considered safe. In addition, EAP-GTC is used in PEAP or TTLS tunnels in wireless environments. The EAP-GTC is described in RFC 2284.
- EAP-AKA—The EAP-AKA authentication mechanism is typically used in mobile networks that include UMTS and CDMA 2000. This method uses the information stored in the SIM for authentication. The EAP-AKA is described in RFC 4187.
- EAP-FAST—The EAP-FAST is an alternative authentication method to PEAP. This method uses the PAC for verifying clients on the network. The EAP-FAST is described in RFC 4851.
- EAP-MD5—The EAP-MD5 method verifies MD5 hash of a user password for authentication. This method is commonly used in a trusted network. The EAP-MD5 is described in RFC 2284.
- EAP-POTP—The EAP type 32 is supported. Complete details are described in RFC 4793.
- EAP-SIM—The EAP-SIM uses GSM 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 uses 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 method allows you to add additional information in an EAP message. Often this method is used to provide more information about an EAP message such as status information or authorization data. This method is always used after a typical EAP authentication process.
- EAP-TTLS—The EAP-TTLS method uses server-side certificates to set up authentication between clients and servers. The actual authentication is, however, performed using passwords. Complete details about EAP-TTLS is described in RFC 5281.
- LEAP— LEAP uses dynamic WEP keys and mutual authentication between the client and the RADIUS server.
- TEAP— TEAP is a tunnel-based EAP method that enables secure communication between a peer and a server by using the Transport Layer Security (TLS) protocol to establish a mutually authenticated tunnel. Further details about TEAP are described in RFC 7170.
- ZLXEAP—ZoneLabs EAP is an EAP method that has been allocated EAP Type 44 by IANA. For more information, visit http://tools.ietf.org/html/draft-bersani-eap-synthesis-sharedkeymethods-00#page-30.

Configuring Authentication with a RADIUS Server

See Table 60 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 34 *802.1X Authentication with a RADIUS Server*



The supplicant and the authentication server must be configured to use the same EAP type. The managed device does not need to know the EAP type used between the supplicant and authentication server.

For the managed device to communicate with the authentication server, you must configure the IP address, authentication port, and accounting port of the server on the managed device. The authentication server must be configured with the IP address of the RADIUS client, which is the managed device in this case. Both the managed device and the authentication server must be configured to use the same shared secret.

NOTE

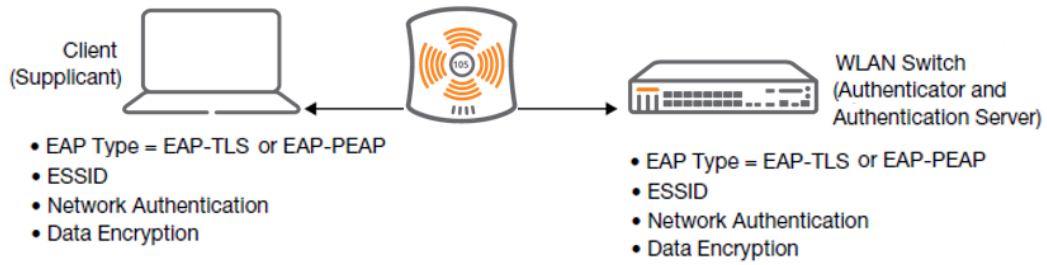
Additional information on EAP types supported in a Windows environment, Microsoft supplicants, and authentication servers, is available at http://technet.microsoft.com/en-us/library/cc782851(WS.10).aspx.

The client communicates with the managed device through a GRE tunnel to form an association with an AP and to get authenticated in the network. Therefore, the network authentication and encryption configured for an ESSID must be the same on both the client and the managed device.

Configuring Authentication Terminated on a Managed Device

User authentication is performed either via the managed device's internal database or a non-802.1X server. See 802.1X Authentication Profile WebUI Parameters on page 272 for an overview of the parameters that you need to configure on 802.1X authentication components when 802.1X authentication is terminated on the managed device (AAA FastConnect).

Figure 35 *802.1X Authentication with Termination on Managed device*



In this scenario, the supplicant is configured for EAP-TLS or EAP-PEAP.

- EAP-TLS is used with smart card user authentication. A smart card holds a digital certificate which, with the user-entered PIN, allows the user to be authenticated on the network. EAP-TLS relies on digital certificates to verify the identities of both the client and the server.
- EAP-TLS requires that you import server and CA certificates onto the managed device (see Configuring 802.1X Authentication on page 271). The client certificate is verified on the managed device (the client certificate must be signed by a known CA) before the username is checked on the authentication server.
- EAP-PEAP uses TLS to create an encrypted tunnel. Within the tunnel, one of the following "inner EAP" methods is used:
  - EAP-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 managed device as a backup to an external authentication server.
  - EAP-Microsoft MS-CHAPv2: Described in RFC 2759, this EAP method is widely supported by Microsoft clients. A RADIUS server must be used as the backend authentication server.

If you use the internal database of the managed device for user authentication, you need to add the names and passwords of the users to be authenticated. If you use an LDAP server for user authentication, you need to configure both the LDAP server and the user IDs and passwords on the managed device. If you use a RADIUS server for user authentication, you need to configure the RADIUS server on the managed device.

Configuring 802.1X Authentication

On the managed device, use the following steps to configure a wireless network that uses 802.1X authentication:

1. Configure the VLANs to which the authenticated users will be assigned. See Network Configuration Parameters on page 80.
2. Configure policies and roles. You can specify a default role for users who are successfully authenticated using 802.1X. You can also configure server derivation rules to assign a user role based on attributes returned by the authentication server; server-derived user roles take precedence over default roles. For more information about policies and roles, see Roles and Policies on page 491. The Policy Enforcement Firewall Virtual Private Network module provides identity-based security for wired and wireless users and must be installed on the managed device. The stateful firewall allows user classification based on user identity, device type, location, and time of day to provide differentiated access for different classes of users. For information about obtaining and installing licenses, refer to the *Aruba Mobility Master Licensing Guide.*
3. Configure the authentication server(s) and server group. The server can be an 802.1X RADIUS server or, if you use AAA FastConnect, a non-802.1X server or the internal database of the managed device. If you use EAP-GTC within a PEAP tunnel, configure an LDAP or RADIUS server as the authentication server (see Authentication Servers on page 192). If you use EAP-TLS, import server and CA certificates on the managed device (see Configuring and Using Certificates with AAA FastConnect on page 277).
4. Configure the AAA profile:
   a. Select the 802.1X default user role.
   b. Select the server group you previously configured for the 802.1X authentication server group.
5. Configure the 802.1X authentication profile. See Example Configurations on page 281.
6. Configure the virtual AP profile for an AP group or for a specific AP:
   a. Select the AAA profile you previously configured.
   b. 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 281.

The following procedure describes how to create and configure a new instance of an 802.1X authentication profile:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > L2 Authentication** tab.
2. In the **L2 Authentication** table, select **802.1X Authentication**.
3. Click **+** in the **802.1X Authentication Profile: New Profile.**
4. Enter a **Profile Name.**
5. Change the settings described in Table 60 as desired**.**
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Table 60: *802.1X Authentication Profile WebUI Parameters*

Parameter	Description
Max Authentication Failures	Number of times a user can try to log in with wrong credentials after which the user is blacklisted as a security threat. Set to 0 to disable blacklisting, otherwise enter a non-zero integer to blacklist the user after the specified number of failures. Range: 0-5 failures. Default: 0 failure.

Table 60: *802.1X Authentication Profile WebUI Parameters*

Parameter	Description
	NOTE: This option may require a license.
Enforce Machine Authentication	Select the Enforce Machine Authentication option to require machine authentication. This option is also available on the **Basic** settings tab. **NOTE:** This option may require a license.
Machine Authentication: Default Machine Role	Default role assigned to the user after 802.1X 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 this check box 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 completing only machine authentication. The default role for this setting is the "guest" role.
Interval between Identity Requests	Interval, in seconds, between identity request retries. Range: 1-65535 seconds. Default: 5 seconds.
Quiet Period after Failed Authentication	The enforced quiet period interval, in seconds, following failed authentication. Range: 1-65535 seconds. Default: 30 seconds.
Reauthentication Interval	Interval, in seconds, between reauthentication attempts. Range: 60-864000 seconds. Default: 86400 seconds (1 day).
Use Server provided Reauthentication Interval	Select this option to override any user-defined reauthentication interval and use the reauthentication period defined by the authentication server.
Use the termination - action attribute from the Server	Select this option to honor termination- action attribute from the server.
Multicast Key Rotation Time Interval	Interval, in seconds, between multicast key rotation. Range: 60-864000 seconds. Default: 1800 seconds.
Unicast Key Rotation Time Interval	Interval, in seconds, between unicast key rotation. Range: 60-864000 seconds. Default: 900 seconds.
Authentication Server Retry Interval	Server group retry interval, in seconds. Range: 2-65535 seconds. Default: 5 seconds.

Table 60: *802.1X Authentication Profile WebUI Parameters*

Parameter	Description
Authentication Server Retry Count	Maximum number of authentication requests that are sent to server group. Range: 0-5 requests. Default: 3 requests.
Framed MTU	Sets the framed MTU attribute sent to the authentication server. Range: 500-1500 bytes. Default: 1100 bytes.
Max number of requests sent during an Auth attempt	Maximum number of times ID requests are sent to the client. Range: 1-10 retries. Default: 5 retries.
Maximum Number of Reauthentication Attempts	Number of times a user can try to log in with wrong credentials after which the user is blacklisted as a security threat. Set to 0 to disable blacklisting, otherwise enter a value from 0-5 to blacklist the user after the specified number of failures. Default: 3 retries. **NOTE:** If changed from its default value, this option may require a license.
Maximum number of times Held State can be bypassed	Number of consecutive authentication failures which, when reached, causes the managed device to not respond to authentication requests from a client while the managed device is in a held state after the authentication failure. Before this number is reached, the managed device responds to authentication requests from the client even while the managed device is in its held state. (This parameter is applicable when 802.1X authentication is terminated on the managed device, 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 or WPA2 key messages are retried. Range: 1-5 retries. Default: 1 retry.
Dynamic WEP Key Size	The default dynamic WEP key size is 128 bits, If desired, you can change this parameter to 40 bits.
Interval between WPA/WPA2 Key Messages	Interval, in milliseconds, between each WPA key exchanges. Range: 1000-5000 ms. Default: 1000 ms.
Delay between EAP-Success and WPA2 Unicast Key Exchange	Interval, in milliseconds, between EAP-Success and unicast key exchanges. Range: 0-2000 ms. Default: 0 ms (no delay).
Delay between WPA/WPA2 Unicast Key and Group Key Exchange	Interval, in milliseconds, between unicast and multicast key exchange. Time interval in milliseconds. Range: 0-2000. Default: 0 (no delay).
Time interval after which the PMKSA will be deleted	The time interval after which the Pairwise Master Key Security Association cache is deleted. Time interval in Hours. Range: 1-2000. Default: 8.

Table 60: *802.1X Authentication Profile WebUI Parameters*

Parameter	Description
WPA/WPA2 Key Message Retry Count	Number of times WPA or WPA2 key messages are retried. Range: 1-5 retries. Default: 3 retries.
Multicast Key Rotation	Select this check box to enable multicast key rotation. This feature is disabled by default.
Unicast Key Rotation	Select this check box to enable unicast key rotation. This feature is disabled by default.
Reauthentication	Select the Reauthentication check box 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 PMK which is derived through a client and an associated AP. This key is used when the client roams to a new AP. This allows clients faster roaming without a full 802.1X authentication. Uncheck this option to disable this feature. **NOTE:** Make sure that the wireless client (the 802.1X supplicant) supports this feature. If the client does not support this feature, the client will attempt to renegotiate the key whenever it roams to a new AP. As a result, the key cached on the managed device can be out of sync with the key of the client.
Validate PMKID	This parameter instructs the managed device to check the PMK ID sent by the client. When you enable this option, the client must send a PMK ID in the associate or reassociate frame to indicate that it supports OKC or PMK caching; otherwise, full 802.1X authentication takes place. **NOTE:** This feature is optional, since most clients that support OKC and PMK caching do not send the PMK ID in their association request.
Use Session Key	Use Session key as the Unicast WEP key. This option is disabled by default.
Use Static Key	Use Static key as Unicast / Multicast WEP key. This option is disabled by default.
xSec MTU	Maximum size used for xSec MTU. Default: 1300
Termination	Select this check box to allow 802.1X authentication to terminate on the managed device. This option is disabled by default.
Termination EAP-Type	If you enable termination, click either EAP-PEAP or EAP-TLS to select a EAP method.
Termination Inner EAP-Type	If you use EAP-PEAP as the EAP method, specify one of the following inner EAP types: ■ **eap-gtc**: Described in RFC 2284, this EAP method permits the transfer of unencrypted usernames and passwords from client to server. The main uses for EAP-GTC are one-time token cards such as SecureID and the use of LDAP

Table 60: *802.1X Authentication Profile WebUI Parameters*

Parameter	Description
	or RADIUS as the user authentication server. You can also enable caching of user credentials on the managed device as a backup to an external authentication server. ■ **eap-mschapv2**: Described in RFC 2759, this EAP method is widely supported by Microsoft clients.
Enforce Suite-B 128 bit or more security level Authentication	Configure Suite-B 128 bit or more security level authentication enforcement.
Enforce Suite-B 128 bit or more security level Authentication	Configure Suite-B 192 bit security level authentication enforcement.
Termination	Select the **Termination** check box to allow 802.1X authentication to terminate on the managed device. This option is disabled by default.
Termination EAP-Type	If you enable termination, click either EAP-PEAP or EAP-TLS to select a EAP method.
Termination Inner EAP-Type	If you use EAP-PEAP as the EAP method, specify one of the following inner EAP types: ■ **eap-gtc**: Described in RFC 2284, this EAP method permits the transfer of unencrypted usernames and passwords from client to server. The main uses for EAP-GTC are one-time token cards such as SecureID and the use of LDAP or RADIUS as the user authentication server. You can also enable caching of user credentials on the managed device 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** check box to enable the managed device to cache the username and password of each authenticated user. The managed device continues to reauthenticate users with the remote authentication server. However, if the authentication server is unavailable, the managed device 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 managed device before it will appear on this list.
Server-Certificate	Click the **Server-Certificate** drop-down list and select a server certificate the managed device will use to authenticate itself to the client. **NOTE:** By default, the **default-self-signed** certificate is used as server certificate. For more details on **default-self-signed** certificate, see Managing Certificates on page 972.
TLS Guest Access	Select **TLS Guest Access** to enable guest access for EAP-TLS users with valid certificates. This option is disabled by default.

Table 60: *802.1X Authentication Profile WebUI Parameters*

Parameter	Description
TLS Guest Role	Click the **TLS Guest Role** drop-down list and select the default user role for EAP-TLS guest users. This option may require a license.
Ignore EAPOL-START after authentication	Select **Ignore EAPOL-START** after authentication to ignore EAPOL-START messages after authentication. This option is disabled by default.
Handle EAPOL-Logoff	Select **Handle EAPOL-Logoff** to enable handling of EAPOL-LOGOFF messages. This option is disabled by default.
Ignore EAP ID during negotiation	Select **Ignore EAP ID during negotiation** to ignore EAP IDs during negotiation. This option is disabled by default.
WPA-Fast-Handover	Select this option to enable WPA-fast-handover on phones that support this feature. WAP fast-handover is disabled by default.
Check certificate common name against AAA server	If you use client certificates for user authentication, enable this option to verify that the common name of the certificate exists in the server. This parameter is enabled by default in the default-cap and default-rap VPN profiles, and disabled by default on all other VPN profiles.

The following CLI command configures settings for an 802.1X authentication profiles. Individual parameters are described in the previous table.

```
(host) [mynode](config)# aaa authentication dot1x {<profile>|countermeasures}
```

Configuring EAP-TLS Fragmentation

The following CLI command configures EAP-TLS fragmentation in an 802.1X authentication profile:

```
(host) [mynode](config) #aaa authentication dot1x eap-frag-mtu <ipmtu>
```

Configuring and Using Certificates with AAA FastConnect

The managed device supports 802.1X authentication using digital certificates for AAA FastConnect.

- Server Certificate—A server certificate installed in the managed device verifies the authenticity of the managed device for 802.1X authentication. Aruba managed device ship with a demonstration digital certificate. Until you install a customer-specific server certificate in the managed device, 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 CA. You can generate a CSR on the managed device to submit to a CA. For information on how to generate a CSR and how to import the CA-signed certificate into the managed device, see Managing Certificates on page 972.
- Client Certificates—Client certificates are verified on the managed device (the client certificate must be signed by a known CA) before the username is checked on the authentication server. To use client certificate authentication for AAA FastConnect, you need to import the following certificates into the managed device (see Managing Certificates on page 972):
  - Server certificate of the managed device
  - CA certificate for the CA that signed the client certificates

The following procedure describes how to configure the server certificate and CA certificate:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > L2 Authentication** tab.
2. Select **802.1X Authentication**.
3. Select the **default** 802.1X authentication profile to display configuration parameters.
4. Select the **Termination** check box.
5. Select the server certificate imported into the managed device from the **Server-Certificate** drop-down list.
6. Select the CA certificate imported into the managed device from the **CA-Certificate** drop-down list.
7. Click **Submit as** and enter a name for the 802.1X authentication profile. Click **Save**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the server certificate and CA certificate:

```
(host)[mynode](config) # aaa authentication dot1x <profile>
      termination enable
      server-cert <certificate>
      ca-cert <certificate>
```

Configuring User and Machine Authentication

When a Windows device boots, it logs onto the network domain using a machine account. Within the domain, the device is authenticated before computer group policies and software settings can be executed; this process is known as *machine authentication*. Machine authentication ensures that only authorized devices are allowed on the network.

You can configure 802.1X for both user and machine authentication (select the **Enforce Machine Authentication** option described in Table 60). This tightens the authentication process further, since both the device and user need to be authenticated.

Working with Role Assignment with Machine Authentication Enabled

When you enable machine authentication, there are two additional roles you can define in the 802.1X authentication profile:

- Machine authentication default machine role
- Machine authentication default user role

While you can select the same role for both options, you should define the roles as per the polices that need to be enforced. Also, these roles can be different from the 802.1X authentication default role configured in the AAA profile.

With machine authentication enabled, the assigned role depends upon the success or failure of the machine and user authentications. In certain cases, the role that is ultimately assigned to a client can also depend upon attributes returned by the authentication server or server derivation rules configured on the managed device.

Table 61 describes role assignment based on the results of the machine and user authentications.

Table 61: *Role Assignment for User and Machine Authentication*

Machine Auth Status	User Auth Status	Description	Role Assigned
Failed	Failed	Both machine authentication and user authentication failed. L2 authentication failed.	No role assigned. No access to the network allowed.
Failed	Passed	Machine authentication failed (for example, the machine information is not present on the server) and user authentication succeeded. Server-derived roles do not apply.	Machine authentication default user role configured in the 802.1X authentication profile.
Passed	Failed	Machine authentication succeeded and user authentication has not been initiated. Server-derived roles do not apply.	Machine authentication default machine role configured in the 802.1X authentication profile.
Passed	Passed	Both machine and user are successfully authenticated. If there are server-derived roles, the role assigned via the derivation take precedence. This is the only case where server-derived roles are applied.	A role derived from the authentication server takes precedence. Otherwise, the 802.1X authentication default role configured in the AAA profile is assigned.

For example, if the following roles are configured:

- 802.1X authentication default role (in AAA profile): dot1x_user
- Machine authentication default machine role (in 802.1X authentication profile): dot1x_mc
- Machine authentication default user role (in 802.1X authentication profile): guest

Role assignment is as follows:

- If both machine and user authentication succeed, the role is dot1x_user. If there is a server-derived role, the server-derived role takes precedence.
- If only machine authentication succeeds, the role is dot1x_mc.
- If only user authentication succeeds, the role is guest.
- On failure of both machine and user authentication, the user does not have access to the network.

With machine authentication enabled, the VLAN to which a client is assigned (and from which the client obtains its IP address) depends upon the success or failure of the machine and user authentications. The VLAN that is ultimately assigned to a client can also depend upon attributes returned by the authentication server or server derivation rules configured on the managed device (see Understanding VLAN Assignments on page 81). If machine authentication is successful, the client is assigned the VLAN configured in the virtual AP profile. However, the client can be assigned a derived VLAN upon successful user authentication.



You can optionally assign a VLAN as part of a user role configuration. Do not use VLAN derivation if you configure user roles with VLAN assignments.

Table 62 describes VLAN assignment based on the results of the machine and user authentications when VLAN derivation is used.

Table 62: *VLAN Assignment for User and Machine Authentication*

Machine Auth Status	User Auth Status	Description	VLAN Assigned
Failed	Failed	Both machine authentication and user authentication failed. L2 authentication failed.	No VLAN.
Failed	Passed	Machine authentication failed (for example, the machine information is not present on the server) and user authentication succeeded.	VLAN configured in the virtual AP profile.
Passed	Failed	Machine authentication succeeded and user authentication has not been initiated.	VLAN configured in the virtual AP profile.
Passed	Passed	Both machine and user are successfully authenticated.	Derived VLAN. Otherwise, VLAN configured in the virtual AP profile.



The administrator can now associate a VLAN ID to a client data based on the authentication credentials in a bridge mode.

Enabling 802.1X Supplicant Support on an AP

ArubaOS provides 802.1X supplicant support on the AP. The AP can be used as a 802.1X supplicant where access to the wired Ethernet network is restricted to those devices that can authenticate using 802.1X. You can provision an AP to act as an 802.1X supplicant and authenticate to the infrastructure using the PEAP protocol. Both Campus APs and Remote APs can be provisioned to use 802.1X authentication.

Prerequisites

- An AP has to be configured with the credentials for 802.1X authentication. These credentials are stored securely in the AP flash.
- The AP must complete the 802.1X authentication before it sends or receives IP traffic such as DHCP.



If the AP cannot complete 802.1X authentication (explicit failure or reply timeout) within 1 minute, the AP will proceed to initiate the IP traffic and attempt to contact the managed device. The infrastructure can be configured to allow this. If the AP contacts the managed device it will be marked as unprovisioned so that the administrator can take corrective action.

Provisioning an AP as an 802.1X Supplicant

The following procedure describes how to provision an AP as an 802.1X supplicant:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > Access Points**.
2. Click **Access Points > Provisioning** window. The list of discovered APs are displayed on this page.
3. Select the AP you want to provision.

4. Select the AP to which you want to add new provisioning settings, then click **Provision**. The AP provisioning settings divided into two groups. By default, the ArubaOS WebUI displays only the basic, commonly used configuration settings. The advanced settings are hidden until you click the **Show Advanced** options link.
5. In the **Uplink authentication** option, select either **EAP-PEAP** or **EAP-TLS** radio button based on your preference
6. (Optional) If you select **EAP-PEAP**, do the following:
   a. Enter the username of the AP in the **User Name** field.
   b. Enter the password of the AP in the **Password** field.
   c. Enter the password again in the **Retype EAP-PEAP password** field and reconfirm it.
7. (Optional) Select **EAP-TLS** radio button. Enable the **Use factory certificates** option if you want the AP to use the factory certificates to perform 802.1x EAP-TLS authentication. When you select **EAP-TLS** with **Use factory certificates** option for AP uplink authentication, you must import Aruba's root CA to the authentication server certificate CA trusted store for TLS client certificate validation.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands provision an AP as an 802.1X supplicant using PEAP:

```
(host) [mynode] (config)#  provision-ap
(host) [mynode] (config-submode)#  apdot1x-username <username>
(host) [mynode] (config-submode)#  apdot1x-passwd <password>
```

The following CLI commands provision an AP as an 802.1X supplicant using EAP-TLS:

```
(host) [mynode] (config)#  provision-ap
(host) [mynode] (config-submode)#  apdot1x-tls
(host) [mynode] (config-submode)#  apdot1x-tls-suffix
(host) [mynode] (config-submode)#  apdot1x-tls-suffix-domain
```

The following CLI command displays the 802.1X authentication details on the managed devices:

```
(host) [mynode] # show ap active
```

If you enable both EAP-PEAP and EAP-TLS methods, the EAP-PEAP authentication takes precedence.



You can add a Fully Qualified Domain Name (FQDN) as a suffix to an AP name or a group of APs for both TPM certificates and EST certificates with EAP-TLS supplicant support. This is done to enable policy differentiation and apply the policies to AP groups in ClearPass or any RADIUS server. Once you enable the **apdot1x-tls-suffix** parameter in the CLI, you can configure the suffix domain and use it as an EAP-TLS username. If you do not configure the suffix domain, the system uses **aruba.ap** as the default domain suffix.

Example Configurations

The following examples show basic configurations:

- Configuring Authentication with an 802.1X RADIUS Server on page 282
- Configuring Authentication with the Internal Database of the Managed Device on page 293
- Configuring Mixed Authentication Modes on page 298

In the following examples:

- Wireless clients associate to the ESSID **WLAN-01**.
- The following roles allow different network's access capabilities:
  - student
  - faculty
  - guest
  - system administrators

Configuring Authentication with an 802.1X RADIUS Server

The examples show how to configure using the WebUI and CLI commands.

- 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 Aruba Mobility Master.
- The authentication type is WPA. From the 802.1X authentication exchange, the client and the Mobility Master 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 users Windows login and password are used for user authentication. A single user sign-on facilitates both authentication to the wireless network and access to the Windows server resources.



802.1X Configuration for IAS and Windows Clients on page 1262 describes how to configure the Microsoft Internet Authentication Server and Windows XP wireless client to operate with the managed device configuration shown in this section.

Configuring Roles and Policies

You can create the following policies and user roles for:

- Student
- Faculty
- Guest
- Sysadmin
- Computer

Creating the Student Role and Policy

The **student** policy prevents students from using telnet, POP3, FTP, SMTP, SNMP, or SSH to the wired portion of the network. The **student** policy is mapped to the **student** user role.

Before creating a student role, it is recommended to create a destination alias **Internal Network.**

The following procedure describes how to create a destination alias.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Roles & Policies** > **Aliases** tab.
2. In **Network Aliases**, click **+**.
3. Select an **IP Version** from the drop-down list.
4. For **Name**, enter **Internalnetwork**

5. For **Description**, enter a description of the destination within 128 characters.
6. Select **Invert** to specify that the inverse of the network addresses configured are used.
7. For **Items**, click **+**.
8. In the **Add New Destination Add New User Rule** window, for Rule Type, select **Network**. For IP Address, enter 10.0.0.0. For Network Mask or Range, enter 255.0.0.0. **Click OK.**
9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following procedure describes how to create a student role:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Policies** tab.
2. Select **+** to add the student policy.
3. For **Policy Name**, enter **student.**
4. For **Policy Type**, select **Session**.
5. Click **Submit**.
6. Select the **student** role from the **Policies** table.
7. Click **+** in the **Policies > student** table to add rules for the policy.
   a. For **Rule type**, select **Access Control**, then click **OK**.
   b. For **Source**, select **User**.
   c. For **Destination**, select **Alias**.
   d. The following step defines an alias representing all internal network addresses. Once defined, you can use the alias for other rules and policies.
   e. For **Destination alias**, select **Internalnetwork**.
   f. For **Service/app**, select **service**. In the Service scrolling list, select **svc-telnet**.
   g. Under **Action**, select **drop**.
   h. Click **Submit**.
8. Repeat step 7 to create rules for the following services: svc-pop3, svc-ftp, svc-smtp, svc-snmp, and svc-ssh.
9. Click **Submit**.
10. Click the **Roles** tab. Click **+** to create the student role.
    a. For Name, enter **student** then click **Submit**.
    b. Select the role you just created from the **Roles** table.
    c. Select **Show Advanced View**.
    d. In the **Roles > student** table, select the **Policies** tab.
    e. Click **+** to add a new policy.
    f. Select **Add existing session policy** and select the student policy you previously created.
11. Click **Submit**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands create a destination alias:

```
(host)[mynode](config) #ip access-list session student
      user alias "Internal Network" svc-telnet deny
      user alias "Internal Network" svc-pop3 deny
      user alias "Internal Network" svc-ftp deny
      user alias "Internal Network" svc-smtp deny
      user alias "Internal Network" svc-snmp deny
      user alias "Internal Network" svc-ssh deny
```

The following CLI command s create the student role and policy:

```
(host)[mynode](config) #user-role student
      session-acl student
      session-acl allowall
```

Creating the Faculty Role and Policy

The **faculty** policy is similar to the **student** policy, however faculty members are allowed to use POP3 and SMTP for VPN remote access from home. (Students are not permitted to use VPN remote access.) The **faculty** policy is mapped to the **faculty** user role.

The following procedure describes how to create the faculty role and policy:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Policies** tab.
2. Click **+** to add the faculty policy.
3. For **Policy Name**, enter **faculty**.
4. For **Policy Type**, select **Session**.
5. Click **Submit**.
6. Select the new **faculty** policy from the Policies table.
7. Click + in the **Policies > Faculty** table to add rules for the policy.
    a. Select the **Rule Type** as **Access Control**, then click **OK**.
    b. For **Source**, and select **User**.
    c. For **Destination**, select **Alias**, then select **Internal Network** for Destination Alias.
    d. For **Service/App**, select **Service**.
    e. For **Service Alias**, select **svc-telnet**.
    f. For **Action**, and select **Deny**.
    g. Click **Submit**.
8. Repeat steps A-E to create rules for the following services: *svc-ftp*, *svc-snmp*, and *svc-ssh*.
9. Select the **Roles** tab. Click + to create the faculty role.
    a. Enter **faculty** for **Name** .
10. Click **Submit**.
11. Select the role you just created from the **Roles** table.
12. Select **Show Advanced View**.
    a. In the **Roles > faculty** table, select the **Policies** tab.
    b. Click **+** to add a new policy.
    c. Select **Add existing session policy** and select the faculty policy you previously created.

13. Click **Submit**.
14. Click **Pending Changes**.
15. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands create the faculty role and policy:

```
(host)[mynode](config) #ip access-list session faculty
      user alias "Internal Network" svc-telnet deny
      user alias "Internal Network" svc-ftp deny
      user alias "Internal Network" svc-snmp deny
      user alias "Internal Network" svc-ssh deny

(host)[mynode](config) #user-role faculty
      session-acl faculty
      session-acl allowall
```

Creating the Guest Role and Policy

The **guest** policy permits only access to the internet (via HTTP or HTTPS) and only during daytime working hours. The **guest** policy is mapped to the **guest** user role.

The following procedure describes how to create the guest role and policy:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies> Roles** tab.
2. Select a role name and click **+** in the **Name of the role > Global roles** table.
3. Click **+** for the **Time range** field and enter the following details:
   a. For Name, enter **working-hours**.
   b. For Type, select **Periodic**. Click **+**.
   c. For **Start day**, click **Weekday**.
   d. For **Start time(hh:mm)**, enter **07:30**.
   e. For **End time(hh:mm)**, enter **17:00**.
4. Click **OK**.
5. Click **Submit**.
6. Click the **Policies** tab. Click **+** to add the guest policy.
   a. For **Policy Name**, enter **guest**.
   b. For **Policy Type**, select **Session**.
7. Click **Submit**.
8. Select the Policy created under **Policies**. The **Policies > policy Name** table is displayed.
9. Click **+** under the **Policies > policy Name** table.
10. Select **Access Control** for the **Rule Type** and click **OK**.
11. Add the following **New Forwarding Rule** information for the policy.
12. To create rules to permit access to DHCP and DNS servers during working hours:
    a. For **Source**, select **User**.
    b. For **Destination**, select **Host**. In Host IP, enter **10.1.1.25**.
    c. For **Service**, select **Service**. In the Service scrolling list, select **svc-dhcp**.

d. For **Action**, select **Permit**.

e. For **Time Range**, select **working-hours**.

13. Click **Submit**.
14. Repeat steps A-F to create a rule for *svc-dns*.
15. To create a rule to deny access to the internal network:
    a. For **Source**, select **User**.
    b. For **Destination**, select **alias**. Select **Internal Network**.
    c. Under **Service**, select **Any**.
    d. Under **Action**, select **Deny**
16. Click **Submit**.
17. To create rules to permit HTTP and HTTPS access during working hours:
    a. For **Source**, select **User**.
    b. For **Destination**, select **Any**.
    c. For **Service/app**, select **Service**. In the Services scrolling list, select **svc-http**.
    d. For **Action**, select **Permit**.
    e. For Time Range, select **working-hours**.
18. Click **Submit**.
19. Repeat steps A-F for the *svc-https* service.
20. To create a rule that denies the user access to all destinations and all services:
    a. For **Source**, select **User**.
    b. For **Destination**, select **Any**.
    c. For **Service/app**, select **Any**.
    d. For **Action**, select **drop**.
21. Click **Submit**.
22. Click the **Roles** tab. Click **+** to create the guest role.
23. For Role Name, enter **guest** and click **Submit.**
24. Under **Firewall Policies**, click **+.** In Choose from Configured Policies, select the guest policy you previously created. Click **Submit.**
25. Click **Pending Changes**.
26. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands create the guest role and policy:

```
   (host)[mynode](config) time-range working-hours periodic
      weekday 07:30 to 17:00

(host)[mynode](config) #ip access-list session guest
      user host 10.1.1.25 svc-dhcp permit time-range working-hours
      user host 10.1.1.25 svc-dns permit time-range working-hours
      user alias "Internal Network" any deny
      user any svc-http permit time-range working-hours
      user any svc-https permit time-range working-hours
      user any any deny
```

```
(host)[mynode](config) #user-role guest
      session-acl guest
```

Creating Roles and Policies for Sysadmin and Computer

The **allowall** policy, a predefined policy, allows unrestricted access to the network. The **allowall** policy is mapped to both the **sysadmin** user role and the **computer** user role.

The following procedure describes how to create roles and policies for sysadmin and computer:

1. In a **Managed Network** node hierarchy, navigate to **Configuration > Roles & Policies> Roles** tab. Click **+** to create the sysadmin role.
2. 2. Enter a Role Name in the **Name** field. Enter **sysadmin** or **computer** for the required role.
3. Select the role created.
4. In the **<Name of the role>** table, click **Show Advanced View**.
5. Under Policies, Click **+**. In **Add Policy**, select the **Add existing policy** and select the predefined **allowall** policy from the **Policy Name** drop-down list.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands create roles and policies for sysadmin and computer:

```
(host)[mynode](config) #user-role sysadmin
   session-acl allowall
(host)[mynode](config) #user-role computer
   session-acl allowall
```

Creating an Alias for the Internal Network

The following CLI commands configure an alias for the internal network:

```
(host)[MyNode](config) #netdestination "Internal Network"
   network 10.0.0.0 255.0.0.0
   network 172.16.0.0 255.255.0.0
```

Configuring the RADIUS Authentication Server

Configure the RADIUS server IAS1, with IP address 10.1.1.21 and shared key. The RADIUS server is configured to sent an attribute called Class to the managed device; the value of this attribute is set to either "student," "faculty," or "sysadmin" to identify the group of users. The managed device uses the literal value of this attribute to determine the role name.

On the managed device, 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.

The following procedure describes how to configure the RADIUS authentication server:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** page.
2. In the **All Servers** list, click **+**.
3. In the New Server window, enter **IAS1** for the server name
   a. Enter **10.1.1.21** for the server **IP address/hostname**.
   b. Set the **Type** to **RADIUS**.
4. Click **Submit**.

5. Select the new server from the All Servers list.
   a. In the **Shared Key** field, enter a key, such as |*a^t%183923!. (You must enter the key string twice.)
6. Click **Submit**.
7. In the **Server Groups** list, click **+**.
8. Enter the server name **IAS** and click **Submit.**
9. Select the server group **IAS** to display configuration parameters for the server group.
   a. In the **Server Group > IAS** table, click **+**.
10. Select **Add existing server**, select IAS1, then click **Submit**.
    a. In the Server Groups table, select the IAS server group. The **Server Group > IAS** table appears.
    b. In the **Server Group > IAS** table, select **Server Rules**.
    c. Click **+** to add a new server rule.
    d. For **Attribute**, select an attribute from the drop-down list.
    e. For **Operation**, select **value-of** from the drop-down list.
    f. For **Action**, select **set role**.
11. Click **Submit**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the RADIUS authentication server:

```
[host][mynode](config) #aaa authentication-server radius IAS1
  host 10.1.1.21
  key |*a^t%183923!

[host][mynode](config) #aaa server-group IAS
  auth-server IAS1
  set role condition Class value-of
```

Configuring 802.1X Authentication

An AAA profile specifies the 802.1X authentication profile and 802.1X server group to be used for authenticating clients for a WLAN. The AAA profile also specifies the default user roles for 802.1X and MAC authentication.

In the 802.1X authentication profile, configure enforcement of machine authentication before user authentication. If a user attempts to log in before machine authentication completes, the user is placed in the limited guest role.

The following procedure describes how to configure 802.1X authentication:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > L2 Authentication** page.
2. Select **802.1X Authentication** Profile.
3. Select the Profile name.

4. Select **Enforce Machine Authentication**.
    a. For the **Machine Authentication: Default Machine Role**, select **computer**.
    b. For the **Machine Authentication: Default User Role**, select **guest**.
5. Click **Submit**
6. In the **Configuration > Authentication > AAA Profiles** tab.
    a. Expand AAA Profiles, click **+ in AAA Profile: New Profile** to add a new profile.
    b. Enter **aaa_dot1x** as the **Profile Name.**
    c. For **MAC Authentication Default Role**, select **computer**.
    d. For **802.1X Authentication Default Role**, select **faculty**.
7. Click **Submit**.
    a. In the **Profiles** list (under the aaa_dot1x profile), select **802.1X Authentication** Profile.
    b. From the drop-down list, select the **dot1x** authentication profile you configured previously.
8. Click **Submit**
    a. In the **Profiles** list (under the aaa_dot1x profile), select **802.1X Authentication Server Group**.
    b. From the drop-down list, select the **IAS** server group you created previously.
9. Click **Submit.**
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure 802.1X authentication:

```
(host)[mynode](config) #aaa authentication dot1x dot1x
   machine-authentication enable
   machine-authentication machine-default-role computer
   machine-authentication user-default-role guest

(host)[mynode](config) #aaa profile aaa_dot1x
   d>ot1x-default-role faculty
   mac-default-role computer
   authentication-dot1x dot1x
   d>ot1x-server-group IAS
```

Configuring VLANs

In this example, wireless clients are assigned to either VLAN 60 or 61 while guest users are assigned to VLAN 63. VLANs 60 and 61 split users into smaller IP subnetworks, improving performance by decreasing broadcast traffic. The VLANs are internal to the Aruba managed device only and do not extend into other parts of the wired network. The clients' default gateway is the Aruba managed device, 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.

The following procedure describes how to configure VLANs:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces > VLANs** page. Click **+** to add **VLAN_60.**
    a. Enter a **VLAN name**.
    a. For **VLAN ID**, enter **60**.

b. Click **Submit.**

c. Repeat steps A and B to add VLANs 61 and 63.

2. To configure IP parameters for the VLANs, navigate to the **Configuration > Interfaces > VLANs** page.

a. Select **VLAN 60**.

b. Under **VLANs > VLAN_60** table, select the VLAN ID, **60**. Click **IPv4**.

c. For IP Address, enter **10.1.60.1**.

d. For Net Mask, enter **255.255.255.0**.

e. Click **Submit**.

3. Similarly, for VLAN 61, navigate to the **Configuration > Interfaces > VLANs** page.

a. Select **VLAN_61**.

b. Under **VLANs > VLAN_61** table, select the VLAN ID, 61. Click **IPv4**.

c. For IP Address, enter **10.1.61.1**.

d. For Net Mask, enter **255.255.255.0**.

e. Click **Submit**.

4. Similarly, for VLAN 63, navigate to the **Configuration > Interfaces > VLANs** page.

a. Select **VLAN_63**.

b. Under **VLANs > VLAN_63** table, select the VLAN ID, 61. Click **IPv4**.

a. For IP Address, enter **10.1.63.1**.

b. For Net Mask, enter **255.255.255.0.**

c. Click **Submit**.

5. Select the **IP Routes** tab.

a. Click **+** in the Static Default Gateway table.

a. For **IP address**, enter **10.1.1.254**.

b. Click **Submit**.

The following CLI commands configure VLANs:

```
(host)[mynode](config) #vlan 60
(host)[mynode](config) #interface vlan 60
   ip address 10.1.60.1 255.255.255.0
   ip helper-address 10.1.1.25

(host)[mynode](config) #vlan 61
(host)[mynode](config) #interface vlan 61
   ip address 10.1.61.1 255.255.255.0
   ip helper-address 10.1.1.25

(host)[mynode](config) #vlan 63
(host)[mynode](config) #interface vlan 63
   ip address 10.1.63.1 255.255.255.0
   ip helper-address 10.1.1.25

(host)[mynode](config) #ip default-gateway 10.1.1.254
```

Configuring the WLANs

In this example, default AP parameters for the entire network are: the default ESSID is WLAN-01 and the encryption mode is TKIP. A second ESSID called “guest” has the encryption mode set to static WEP with a configured WEP key.

In this example, the non-guest clients that associate to an AP are mapped into one of two different user VLANs. The initial AP to which the client associates determines the VLAN: clients that associate to APs in the first floor of the building are mapped to VLAN 60, and clients that associate to APs in the second floor of the building are mapped to VLAN 61. Therefore, the APs in the network are segregated into two AP groups, named first-floor and second-floor. (See AP Groups on page 647 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.

The following procedure describes how to configure guest WLAN:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Under **All Profiles**, select **Wireless LAN** and then select **SSID**
3. Click **+** in the **SSID Profile: New Profile**. Enter the **Profile Name** and **ESSID** as **guest.**
4. For **Encryption**, select **static-wep.** Click **Submit.**
5. Select **Virtual AP** under **Wireless LAN**.
6. Click **+** in the **Virtual AP Profile: New Profile**.
7. Enter the **Profile Name** as **guest** and select **Virtual AP enable**. Enter a value for **VLAN**.
8. Click **Submit**.
9. Select the Virtual AP created and select **SSID**. Select **guest** from the **SSID profile** drop-down list. Click **Submit.**
10. Navigate to **Configuration > AP groups**.
11. In the **AP Groups** list, select an AP group. In the **APgroups > <name of the group>** table, click the **WLANs** tab.
12. Click **+**.
13. Select the Virtual AP **guest** from the Virtual - AP drop-down list. Click **Submit.**
14. Click **Submit**.
15. Click **Pending Changes**.
16. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure guest WLAN:

```
(host)v(config) #wlan ssid-profile guest
   essid guest
   wepkey1 aaaaaaaaaa
   opmode static-wep

(host)[mynode](config) #wlan virtual-ap guest
   vlan 63
   ssid-profile guest

(host)[mynode](config) #ap-group first-floor
   virtual-ap guest
(host)(config) #ap-group second-floor
   virtual-ap guest
```

Configuring the Non-Guest WLANs

You create and configure the SSID profile “WLAN-01” with the ESSID “WLAN-01” and WPA TKIP encryption. You need to create and configure two virtual AP profiles: one with VLAN 60 for the first-floor AP group and the other with VLAN 61 for the second-floor AP group. Each virtual AP profile references the SSID profile “WLAN-01” and the previously-configured AAA profile aaa_dot1x.

The following procedure describes how to configure the non-guest WLANs:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Under **All Profiles**, select **Wireless LAN** and then select **SSID**
3. Click **+** in the **SSID Profile: New Profile**. Enter the **Profile Name** and **ESSID** as **WLAN-01.**
4. For **Encryption**, select **wpa-tkip.** Click **Submit.**
5. Select **Virtual AP** under **Wireless LAN**.
6. Click **+** in the **Virtual AP Profile: New Profile**.
7. Enter a **Profile Name** for the first floor AP and select **Virtual AP enable**. Enter **60** for **VLAN**.
8. Click **Submit**.
9. Select the Virtual AP created and select **SSID**. Select **WLAN-01** from the **SSID profile** drop-down list. Click **Submit.**
10. Select the Virtual AP created and select **AAA** profile. Select the previously-configured **aaa_dot1x** profile from the **AAA drop down-list.**
11. Repeat steps 5 to 10 to create and associate SSID and AAA profiles for the second floor Virtual AP. Enter the **VLAN** as **61** for the second floor Virtual AP.
12. Navigate to **Configuration > AP groups**.
13. In the **AP Groups** list, select **first-floor**. In the **APgroups > first-floor** table, click the **WLANs** tab.
14. Click **+**.
15. Select the Virtual AP created for first-floor from the Virtual-Ap drop-down list. Click **Submit.**
16. Repeat steps 13 and 14 and select the Virtual AP created for second-floor from the Virtual-Ap drop-down list.
17. Click **Submit**.
18. Click **Pending Changes**.
19. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the non-guest WLANs:

```
(host)[mynode](config) #wlan ssid-profile WLAN-01
   essid WLAN-01
   opmode wpa-tkip

(host)[mynode](config) #wlan virtual-ap WLAN-01_first-floor
   vlan 60
   aaa-profile aaa_dot1x
   ssid-profile WLAN-01

(host)[mynode](config) #wlan virtual-ap WLAN-01_second-floor
   vlan 61
   aaa-profile aaa_dot1x
   sid-profile WLAN-01

(host)[mynode](config) #ap-group first-floor
   virtual-ap WLAN-01_first-floor
(host)[mynode](config) #ap-group second-floor
   virtual-ap WLAN-01_second-floor
(host)[mynode](config) #wlan ssid-profile WLAN-01
   essid WLAN-01
   opmode wpa-tkip

(host)[mynode](config) #wlan virtual-ap WLAN-01_first-floor
   vlan 60
   aaa-profile aaa_dot1x
   ssid-profile WLAN-01
```

```
(host)[mynode](config) #wlan virtual-ap WLAN-01_second-floor
   vlan 61
   aaa-profile aaa_dot1x
   ssid-profile WLAN-01

(host)[mynode](config) #ap-group first-floor
   virtual-ap WLAN-01_first-floor
   ap-group second-floor
   virtual-ap WLAN-01_second-floor
```

Configuring Authentication with the Internal Database of the Managed Device

In the following example:

- The internal database of the managed device provides user authentication.
- The authentication type is WPA. From the 802.1X authentication exchange, the client and the managed device 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.

The following procedure describes how to configure the internal database:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** tab.
2. In the **All Servers** list, select **Internal**.
3. Select a server name under the **Server > Internal** table or Click **+** to add a new server. User name can be entered only for a new server. The User name for an already existing server cannot be changed.
4. For each user, enter a **Password**.
5. Select a **Role** for each user (if a role is not specified, the default role is guest).
6. Select the **Expiration time** for the user account in the internal database.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures the internal database:

```
(host)[mynode](config) #local-userdb add username <user> password <password>
```



Use the privileged mode in the CLI to configure users in the internal database of the managed device.

Configuring a Server Rule

The following procedure describes how to configure a server rules:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** page.
2. Select the **internal** server group from the **Server Groups** table.
3. Click **Server Rules** tab in the **Server Group > Internal** table.

4. Click **+** to add a server derivation rule.
   a. Select an **Attribute.**
   b. Select **value-of** from the **Operations** drop-down list.
   c. Select **Set Role** from the **Action** drop-down list.
   d. Click **Add**.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure server rules:

```
(host)[mynode](config) #aaa server-group internal
   set role condition Role value-of
```

Configuring 802.1X Authentication

An AAA profile specifies the 802.1X authentication profile and 802.1X server group to be used for authenticating clients for a WLAN. The AAA profile also specifies the default user role for 802.1X authentication.

For this example, you enable both 802.1X authentication and termination on the managed device.

The following procedure describes how to configure 802.1X authentication:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > L2 Authentication** tab. In the profiles list, select **802.1X Authentication** Profile.
2. Click **+** in **802.1x Authentication: New Profile.**
   a. For **Profile Name**, enter **dot1x.**
   b. Select **Termination** check box.
   c. Click **Submit**.



The defaults for EAP Method and Inner EAP Method are EAP-PEAP and EAP-MSCHAPv2, respectively.

3. Select the **AAA Profiles** tab and expand **AAA Profiles**.
   a. In the **AAA Profile: New Profile**, click **+** to add a new profile.
   b. Enter **aaa_dot1x** for **Profile Name.**
   c. For 802.1X Authentication Default Role, select **faculty**.
   d. Click **Submit.**
4. In the **Profiles** list (under the aaa_dot1x profile you just created), select **802.1X Authentication Profile**.
   a. Select the dot1x profile from the **802.1X Authentication Profile** drop-down list.
   b. Click **Submit.**
5. 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 **Submit.**

The following CLI commands configure 802.1X authentication:

```
(host)[mynode](config) #aaa authentication dot1x dot1x
  termination enable

(host)[mynode](config) #aaa profile aaa_dot1x
  d>ot1x-default-role student
  authentication-dot1x dot1x
  d>ot1x-server-group internal
```

Configuring VLANs

In this example, wireless clients are assigned to either VLAN 60 or 61 while guest users are assigned to VLAN 63. VLANs 60 and 61 split users into smaller IP subnetworks, improving performance by decreasing broadcast traffic. The VLANs are internal to the Aruba managed device only and do not extend into other parts of the wired network. The default gateway of the client is the Aruba managed device, 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.

The following procedure describes how to configure VLANs:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces > VLANs** page. Click **+** to add VLAN_60.
   a. Enter a **VLAN name.**
   b. For **VLAN ID**, enter **60**.
   c. Click **Submit**.
   d. Repeat steps A and B to add VLANs 61 and 63.
2. To configure IP parameters for the VLANs, navigate to the **Configuration > Interfaces > VLANs** page.
   a. Select **VLAN_60**.
   b. Under **VLANs > VLAN_60** table, select the VLAN ID, 60. Click **IPv4**.
   c. For **IP address**, enter **10.1.60.1**.
   d. For **Net Mask**, enter **255.255.255.0**.
   e. Click **Submit.**
3. To configure IP parameters for the VLANs, navigate to the **Configuration > Interfaces > VLANs** page.
   a. Select **VLAN_61**.
   b. Under **VLANs > VLAN_61** table, select the VLAN ID, 61. Click **IPv4**.
   c. For IP Address, enter **10.1.61.1**.
   d. For Net Mask, enter **255.255.255.0**.
   e. Click **Submit**.

4. To configure IP parameters for the VLANs, navigate to the **Configuration > Interfaces > VLANs** page.
   a. Select **VLAN_63**.
   b. Under **VLANs > VLAN_63** table, select the VLAN ID, 63. Click **IPv4**.
   c. For IP Address, enter **10.1.63.1**.
   d. For Net Mask, enter **255.255.255.0**.
   e. Click **Submit**.
5. Select the **IP Routes** tab.
   a. Click **+** in the **Static Default Gateway** table.
   b. For **IP address**, enter **10.1.1.254**.
   c. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure VLANs:

```
(host)[mynode](config) #vlan 60
(host)[mynode](config) #interface vlan 60
   ip address 10.1.60.1 255.255.255.0
   ip helper-address 10.1.1.25

(host)[mynode](config) #vlan 61
(host)[mynode](config) #interface vlan 61
   ip address 10.1.61.1 255.255.255.0
   ip helper-address 10.1.1.25

(host)[mynode](config) #vlan 63
(host)[mynode](config) #interface vlan 63
   ip address 10.1.63.1 255.255.255.0
   ip helper-address 10.1.1.25

(host)[mynode](config) #ip default-gateway 10.1.1.254
```

Configuring WLANs

In this example, default AP parameters for the entire network are as follows: the default ESSID is WLAN-01 and the encryption mode is TKIP. A second ESSID called guest has the encryption mode set to static WEP with a configured WEP key.

In this example, the non-guest clients that associate to an AP are mapped into one of two different user VLANs. The initial AP to which the client associates determines the VLAN: clients that associate to APs in the first floor of the building are mapped to VLAN 60, and clients that associate to APs in the second floor of the building are mapped to VLAN 61. Therefore, the APs in the network are segregated into two AP groups, named first-floor and second-floor. (See AP Groups on page 647 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.

The following procedure describes how to configure guest WLAN:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Under **All Profiles**, select **Wireless LAN** and then select **SSID**
3. Click **+** in the **SSID Profile: New Profile**. Enter the **Profile Name** and **ESSID** as **guest.**
4. For **Encryption**, select **static-wep.** Click **Submit.**
5. Select **Virtual AP** under **Wireless LAN**.
6. Click **+** in the **Virtual AP Profile: New Profile**.
7. Enter the **Profile Name** as **guest** and select **Virtual AP enable**. Enter a value for **VLAN**.
8. Click **Submit**.
9. Select the Virtual AP created and select **SSID**. Select **guest** from the **SSID profile** drop-down list. Click **Submit.**
10. Navigate to **Configuration > AP groups**.
11. In the **AP Groups** list, select an AP group. In the **APgroups > <name of the group>** table, click the **WLANs** tab.
12. Click **+**.
13. Select the Virtual AP **guest** from the Virtual - AP drop-down list. Click **Submit.**
14. Click **Submit**.
15. Click **Pending Changes**.
16. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure guest WLAN:

```
(host)[mynode](config) #wlan ssid-profile WLAN-01
  essid WLAN-01
  opmode wpa-tkip

(host)[mynode](config) #wlan virtual-ap WLAN-01_first-floor
  vlan 60
  aaa-profile aaa_dot1x
  ssid-profile WLAN-01

(host)[mynode](config) #wlan virtual-ap WLAN-01_second-floor
  vlan 61
  aaa-profile aaa_dot1x
  sid-profile WLAN-01

(host)[mynode](config) #ap-group first-floor
  virtual-ap WLAN-01_first-floor
(host)[mynode](config) #ap-group second-floor
  virtual-ap WLAN-01_second-floor

(host)[mynode](config) #wlan ssid-profile guest
  essid guest
  wepkey1 aaaaaaaaaa
  opmode static-wep

(host)[mynode](config) #wlan virtual-ap guest
  vlan 63
  ssid-profile guest

(host)[mynode](config) #ap-group first-floor
  virtual-ap guest
(host)[mynode](config) #ap-group second-floor
  virtual-ap guest
```

Configuring the Non-Guest WLANs

You create and configure the SSID profile “WLAN-01” with the ESSID “WLAN-01” and WPA TKIP encryption. You need to create and configure two virtual AP profiles: one with VLAN 60 for the first-floor AP group and the other with VLAN 61 for the second-floor AP group. Each virtual AP profile references the SSID profile “WLAN-01” and the previously-configured AAA profile “aaa_dot1x”.

The following procedure describes how to configure non-guest WLANs:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Under **All Profiles**, select **Wireless LAN** and then select **SSID**
3. Click **+** in the **SSID Profile: New Profile**. Enter the **Profile Name** and **ESSID** as **WLAN-01.**
4. For **Encryption**, select **wpa-tkip.** Click **Submit.**
5. Select **Virtual AP** under **Wireless LAN**.
6. Click **+** in the **Virtual AP Profile: New Profile**.
7. Enter a **Profile Name** for the first floor AP and select **Virtual AP enable**. Enter **60** for **VLAN**.
8. Click **Submit**.
9. Select the Virtual AP created and select **SSID**. Select **WLAN-01** from the **SSID profile** drop-down list. Click **Submit.**
10. Select the Virtual AP created and select **AAA** profile. Select the previously-configured **aaa_dot1x** profile from the **AAA drop down-list.**
11. Repeat steps 5 to 10 to create and associate SSID and AAA profiles for the second floor Virtual AP. Enter the **VLAN** as **61** for the second floor Virtual AP.
12. Navigate to **Configuration > AP groups**.
13. In the **AP Groups** list, select **first-floor**. In the **APgroups > first-floor** table, click the **WLANs** tab.
14. Click **+**.
15. Select the Virtual AP created for first-floor from the Virtual-Ap drop-down list. Click **Submit.**
16. Repeat steps 13 and 14 and select the Virtual AP created for second-floor from the Virtual-Ap drop-down list.
17. Click **Submit**.
18. Click **Pending Changes**.
19. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure non-guest WLANs:

```
(host)[mynode](config) #wlan ssid-profile WLAN-01
   essid WLAN-01
   opmode wpa-tkip

(host)[mynode](config) #wlan virtual-ap WLAN-01_first-floor
   vlan 60
   aaa-profile aaa_dot1x
   ssid-profile WLAN-01

(host)[mynode](config) #wlan virtual-ap WLAN-01_second-floor
   vlan 61
   aaa-profile aaa_dot1x
   sid-profile WLAN-01

(host)[mynode](config) #ap-group first-floor
   virtual-ap WLAN-01_first-floor
(host)[mynode](config) #ap-group second-floor
   virtual-ap WLAN-01_second-floor
```

Configuring Mixed Authentication Modes

Use `l2-auth-fail-through` command to perform mixed authentication which includes both MAC and 802.1X authentication. When MAC authentication fails, enable the `l2-auth-fail-through` command to perform 802.1X authentication.



By default the l2-auth-fail-through command is disabled.

Table 63: *Mixed Authentication Modes*

Authentication	1	2	3	4	5	6
MAC authentication	Success	Success	Success	Fail	Fail	Fail
802.1X authentication	Success	Fail	—	Success	Fail	—
Association	dynamic-wep	No Association	static-wep	dynamic-wep	No Association	static-wep
Role Assignment	802.1X	—	MAC	802.1X	—	logon

Table 63 describes the different authentication possibilities

The following CLI commands configure mixed authentication:

```
(host) [mynode] (config) #aaa profile test
    l2-auth-fail-through
```

Performing Advanced Configuration Options for 802.1X

This section describes advanced configuration options for 802.1X authentication.

Configuring Reauthentication with Unicast Key Rotation

When enabled, unicast and multicast keys are updated after each reauthorization. It is a best practice to configure the time intervals for reauthentication, multicast key rotation, and unicast key rotation to be at least 15 minutes. Ensure that these intervals are mutually prime, and the factor of the unicast key rotation interval and the multicast key rotation interval is less than the reauthentication interval. Unicast key rotation depends upon both the AP or managed device 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

The following procedure describes how to configure re-authentication with unicast key rotation:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > L2 Authentication** tab.
2. Select **802.1X Authentication** and select the name of the profile you want to configure.

3. Enter the following values:
    a. Reauthentication Interval: 6011
    b. Multicast Key Rotation Time Interval: 1867
    c. Unicast Key Rotation Time Interval: 1021
    d. Multicast Key Rotation: (select)
    e. Unicast Key Rotation: (select)
    f. Reauthentication: (select)
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure re-authentication with unicast key rotation:

```
(host) [mynode] (config) #aaa authentication dot1x profile
  reauthentication
  timer reauth-period 6011
  unicast-keyrotation
  timer ukey-rotation-period 1021
  multicast-keyrotation
  timer mkey-rotation-period 1867
```

Application Single Sign-On Using L2 Authentication

This feature allows SSO for different web-based applications using Layer 2 authentication information. SSO for web-based application uses SAML, which happens between the web service provider and an identity provider that the web server trusts. A request made from the client to a web server is redirected to the IDP for authentication. If the user has already been authenticated using L2 credentials, the IDP server already knows the authentication details and returns a SAML response, redirecting the client browser to the web-based application. The user enters the web-based application without needing to enter the credentials again.

Enabling application SSO using L2 network information requires configuration on the managed device and on the IDP server. The Aruba ClearPass Policy Manager is the only IDP supported. The managed device has been optimized to work with ClearPass Policy Manager to provide better functionality as an IDP.

Important Points to Remember

- ClearPass Policy Manager is the only supported IDP.
- SSO occurs after 802.1X authentication. Therefore, SSO after captive portal authentication is not supported. Roles for captive portal and SSO are mutually exclusive and, therefore, a user in the captive portal role cannot perform SSO and vice-versa.
- SSO with VIA is not supported.
- There is a limit on the number of concurrent sessions that can be serviced at a given instant. This limit is set at the webserver level using the **web-server profile web-max-clients** command. The default value is 320 for 7000 Series and 7200 Series managed device platforms and 25 for other managed device platforms. The maximum number of concurrent SSO sessions that can be handled is dependent on the other web services being handled and the same time.

Enabling Application SSO

Enabling application SSO using L2 authentication information requires configuration on the managed device and ClearPass Policy Manager. This feature is enabled by completing the following steps:

- ClearPass Policy Manager (refer to the ClearPass Policy Manager for configuration of the following procedures):
  - Add the IP address of the managed device as a network device
  - Add the user to the local user DB
  - Create an enforcement profile to return the Aruba VSA SSO token
  - Create an IDP attribute enforcement profile
  - Create an enforcement policy binding the Aruba VSA SSO token enforcement profile
  - Create an enforcement policy binding the IDP enforcement profile
  - Create a service, allowing the respective authentication types and authentication database, and bind the Aruba VSA SSO token enforcement policy.
  - Create a service, allowing the respective authentication types and authentication database, and bind the IDP enforcement policy.
  - Configure SSO for the ClearPass Policy Manager.
- Managed device:
  - Configuring an SSO-IDP Profile
  - Applying an SSO Profile to a User Role
  - Selecting an IDP Certificate

Configuring SSO IDP-Profiles on the Managed devices

Before SSO can be enabled, you must configure an SSO profile by completing the procedure detailed below.

The following procedure describes how to configure the SSO IDP-Profiles:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles.**
2. Expand **Wireless LAN**.
3. Select **SSO**.
4. In **SSO Profile: New Profile**, click **+** to create a new profile.
5. Enter a **Profile Name**.
6. In **SSO Profiles**, click **+** to add a new URL.
7. Enter the **URL Name**.
8. Enter the **URL**.
9. Click **OK**.
10. Repeat steps 4 through 8 for each URL you are adding to the SSO profile.
11. Click **Submit**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the SSO IDP-Profiles:

```
(host) [mynode] (config) #sso idp-profile <idp profile name>
#idp <urlname> <url>
```

Applying an SSO Profile to a User Role

The newly created SSO profile must be applied to any applicable user rules that require SSO. The following procedure describes how to apply the SSO profile to a user role:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Roles** tab.
2. In the **Roles** table, select the User Role that the SSO profile will be linked to.
3. Click **Show Advanced View**
4. Click the **More** Tab.
5. Expand the **Authentication** accordion.
6. Select an IDP profile from **IDP profile** drop-down list
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands apply an SSO profile to a user role:

```
(host) [mynode](config)# user-role <role name>
(host) [mynode](config)# sso <idp profile name>
```

Selecting an IDP Certificate

An SSL certificate is needed for SSL negotiation with browser. The certificate can be imported in PKCS12 format, so that it contains the certificate and private key, or the key pair can be generated and a CSR request sent to the enterprise CA server to generate a certificate which can then be uploaded to the managed device.

For information about uploading or generating a certificate, see Managing Certificates.

After a certificate is uploaded or generated, the IDP certificate must be selected.

The following procedure describes how to select an IDP certificate:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > System > More** tab.
2. Expand the **General** accordion.
3. Select the IDP certificate from the **IDP Server Certificate** drop-down list. By default, the **default-self-signed** certificate is used as the server certificate. For more details on **default-self-signed** certificate, see Managing Certificates on page 972.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands select an IDP certificate:

```
(host) [mynode](config)# web-server profile
(host) [mynode](Web Server Configuration) # idp-cert <name of the certificate>
```

Device Name as User Name for Non-802.1X Authentication

When a client is authenticated by non-802.1X method of authentication, the host name of the host device is used as the user name (instead of the MAC address) of the host device. When a device tries to obtain an IP address by using DHCP, the host name of the host device in the option-12 field of DHCP request is used as the host name of the device.

A CLI command allows the use of host name or MAC address of a device as the user name of the host device. By default, the MAC address of the host device is used as the user name. If the CLI command is enabled, the host name of the host device is used as the user name.

Using Device Name as User Name

The following CLI commands configure user name:

```
(host) [mm] (config) #aaa profile <profile>
(host) [mm] (AAA Profile "<profile >") #username-from-dhcp-opt12
```

Enhanced Open Security

Enhanced open replaces open unencrypted wireless networks thereby mitigating exposure of user data to passive traffic sniffing. With enhanced open, the client and WLAN perform Diffie-Hellman key exchange during the access procedure and use the resulting pairwise key with a 4-way handshake. ArubaOS supports:

- Enhanced open without PMK caching
- Enhanced open with PMK caching
- enhanced open transition mode

Enhanced Open without PMK Caching

In enhanced open without PMK caching, the 802.11 beacon, probe response frame, and authentication request or response frame are generic. However, the 802.11 association request or response are specific for enhanced open without PMK caching.

ArubaOS advertises support for enhanced open by using an AKM suite selector in all beacons and probe response frames. Besides, PMF is set to required (MFPR=1). Authentication request and authentication response use open authentication.

A client that wishes to perform data encryption in an open Wi-Fi network using enhanced open, indicates enhanced open AKM in the 802.11 association request with PMF is required (MFPR=1). The DHPE contains group and the Diffie-Hellman public Key from the client. ArubaOS supports Diffie-Hellman Groups 19, 20, and 21.

ArubaOS includes the enhanced open AKM and DHPE in the 802.11 association response after agreeing to enhanced open with PME is required (MFPR=1). The DHPE contains group and the Diffie-Hellman public key from ArubaOS. If ArubaOS does not support the group indicated in the received 802.11 association request, it responds with an 802.11 association response having the status code 77. A status code 77 indicates unsupported finite cyclic group.

After completing the 802.11 association, PMK and its associated PMKID are created. ArubaOS initiates a 4-way handshake with the client using the generated PMK. The result of the 4-way handshake is the encryption key to protect bulk unicast data and broadcast data between the client and ArubaOS.

Enhanced Open with PMK Caching

If enhanced open has been established earlier, a client that wishes to perform enhanced open with PMK caching includes a PMKID in its 802.11 association request in addition to the enhanced open AKM, DHPE, and PMF is required(MFPR=1). If ArubaOS has cached the PMK identified by that PMKID, it includes the PMKID in its 802.11 association response but does not include the DHPE. If ArubaOS has not cached the PMK identified by that PMKID, it ignores the PMKID and proceeds with enhanced open association by including a DHPE. The 4-way handshake is initiated subsequently.

Enhanced Open Transition Mode

The enhanced open transition mode enables a seamless transition from open unencrypted WLAN connections without adversely impacting the end user experience. It provides the ability for enhanced open and non-enhanced open clients to connect to the same open system virtual AP.

Two different SSIDs are created for each configured 802.11 open system virtual AP, one for enhanced open and another for open networks. Both SSIDs operate either in the same band and channel or the band and channel of the other SSID (the enhanced open transition mode information element includes the band and channel information). ArubaOS always uses the same band and channel.

802.11 beacon and probe response frames of the open BSS include an enhanced open transition mode information element to encapsulate BSSID and SSID of the enhanced open BSS.

802.11 beacon and probe response frames from the enhanced open BSS include an enhanced open transition mode information element to encapsulate the BSSID and SSID of the open BSS. Besides, the beacon frame from the enhanced open BSS has zero length SSID and indicates enhanced open in robust security network element.

In enhanced open transition mode, ArubaOS uses more virtual APs than configured. The number of virtual APs pushed depends on MultiZone parameters, if configured (maximum SSIDs per zone). During enhanced open transition mode, depending on the available VAP slots, ArubaOS will either push both open and enhanced open virtual APs or only enhanced open virtual APs. There will be no impact on other virtual APs configured. An additional enhanced open virtual AP will be pushed to an AP only if it has an available extra slot.

During transition, if there are many enhanced open enabled virtual APs, based on the availability of slots, the AP will choose to transition all enhanced open virtual APs or configure them as enhanced open-only virtual APs. That is, if there are 2 enhanced open virtual APs and 4 available slots, the AP will create 2 enhanced open-only virtual APs and 2 open virtual APs. If the available slots are 3, the AP will create 2 enhanced open-only virtual APs and no open virtual APs.

Configuring Enhanced Open

The following CLI commands enable enhanced open:

```
(host) [mynode] #configure terminal
(host) [mynode] (config) #wlan ssid-profile enhanced_open_mode
(host) [mynode] (SSID Profile "enhanced_open_mode") #opmode enhanced-open
```

The following procedure describes how to enable enhanced open:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN > SSID**.
3. To create a new SSID profile, click **+** and enter a name for the new SSID profile in **Profile name**.
4. Configure your SSID settings. The configuration parameters are described in WLAN SSID Profiles.
5. In **Encryption**, select **enhanced-open**.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands disable enhanced open:

```
(host) [mynode] #configure terminal
(host) [mynode] (config) #wlan ssid-profile enhanced_open_mode
(host) [mynode] (SSID Profile "enhanced_open_mode") #no opmode
```

The following procedure describes how to disable enhanced open:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN > SSID**.
3. To create a new SSID profile, click **+** and enter a name for the new SSID profile in **Profile name**.

4. Configure your SSID settings. The configuration parameters are described in WLAN SSID Profiles.
5. In **Encryption**, unselect **enhanced-open**.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands enable opmode transition:

```
(host) [mynode] #configure terminal
(host) [mynode] (config) #wlan ssid-profile enhanced_open_mode
(host) [mynode] (SSID Profile "enhanced_open_mode") #opmode-transition
```

The following procedure describes how to enable opmode transition:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN > SSID**.
3. To create a new SSID profile, click **+** and enter a name for the new SSID profile in **Profile name**.
4. Configure your SSID settings. The configuration parameters are described in WLAN SSID Profiles.
5. Select **Opmode transition**.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands disable opmode transition:

```
(host) [mynode] #configure terminal
(host) [mynode] (config) #wlan ssid-profile enhanced_open_mode
(host) [mynode] (SSID Profile "enhanced_open_mode") #no opmode-transition
```

The following procedure describes how to disable opmode transition:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN > SSID**.
3. To create a new SSID profile, click **+** and enter a name for the new SSID profile in **Profile name**.
4. Configure your SSID settings. The configuration parameters are described in WLAN SSID Profiles.
5. Unselect **Opmode transition**.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands display the enhanced open transition mode virtual APs:

```
(host) [mynode] #show ap owe-tm-wins ap-name <ap-name>
(host) [mynode] #show ap owe-tm-wins ip-addr <ip-addr>
(host) [mynode] #show ap owe-tm-wins ip6-addr <ip6-addr>
(host) [mynode] #show ap owe-tm-wins wired-mac <wired-mac>
```

The following CLI commands display the virtual APs that are rejected during enhanced open transition:

```
(host) [mynode] #show ap details advanced ap-name <ap-name>
(host) [mynode] #show ap details advanced ip-addr <ip-addr>
(host) [mynode] #show ap details advanced ip6-addr <ip6-addr>
(host) [mynode] #show ap details advanced wired-mac <wired-mac>
```

Enhanced Open in Decrypt-Tunnel Mode

ArubaOS supports enhanced open in decrypt-tunnel mode.

The following CLI commands configure enhanced open in decrypt-tunnel mode:

```
(host) [mynode] #configure terminal
(host) [mynode] (config) #wlan virtual-ap enhanced_open_mode
(host) [mynode] (Virtual AP profile "enhanced_open_mode") #forward-mode decrypt-tunnel
(host) [mynode] (Virtual AP profile "enhanced_open_mode") #wlan ssid-profile enhanced_
open_test
(host) [mynode] (SSID Profile "enhanced_open_test") #opmode enhanced-open
```

Support for WPA3

ArubaOS supports new WPA3 security improvements with the following features:

- Simultaneous Authentication of Equals (SAE) replaces WPA2-PSK with a password based authentication resistant to dictionary attacks.
- WPA3-Enterprise optionally adds usage of Suite-B 192-bit minimum-level security suite aligned with Commercial National Security Algorithm (CNSA) for enterprise networks. SAE-based keys are not based on PSK and are therefore pairwise and unique between clients and the AP. Suite B restricts the deployment to one of two options:
  - 128-bit security
  - 192-bit security without the ability to mix-and-match ciphers, Diffie-Hellman groups, hash functions, and signature modes

SAE

SAE replaces the less-secure WPA2-PSK authentication. Instead of using the PSK as the PMK, SAE arrives at a PMK, by mapping the PSK to an element of a finite cyclic group, PassWord Element (PWE), doing FCG operations on it, and exchanging it with the peer. ArubaOS supports:

- SAE without PMK caching
- SAE with PMK caching
- SAE or WPA2-PSK mixed mode

SAE Without PMK Caching

ArubaOS advertises support for SAE by using an AKM suite selector for SAE in all beacons and probe response frames. Besides, PMF is set to required (MFPR=1).

A client that wishes to perform SAE sends an 802.11 authentication request with authentication algorithm set to value 3 (SAE). This frame contains a well-formed commit message, that is, authentication transaction sequence set to 1, an FCG, commit-scalar, and commit-element.

ArubaOS responds with an 802.11 authentication containing its own commit message.

ArubaOS and the client compute the PMK and send the confirm message to each other using an authentication frame with authentication transaction sequence set to 2.

The client sends an association request with the AKM suite set to SAE and ArubaOS sends an association response.

ArubaOS initiates a 4-way key handshake with the client to derive the PTK.

SAE With PMK Caching

If SAE has been established earlier, a client that wishes to perform SAE with PMK caching sends an authentication frame with authentication algorithm set to open. ArubaOS sends an authentication response and the client sends a reassociation request with AKM set to SAE and includes the previously derived PMKID.

ArubaOS checks if the PMKID is valid and sends an association response with the status code success.

ArubaOS initiates a 4-way key handshake with the client to derive the PTK.

SAE or WPA2-PSK Mixed Mode

SAE or WPA2-PSK mixed mode allows both SAE clients and clients that can only perform WPA2-PSK to connect to the same BSSID. In this mode, the beacon or probe responses contain a AKM list which contains both PSK (00-0F-AC:2) and SAE (00-0F-AC:8). Client that support SAE send an authentication frame with SAE payload and connect to the BSSID.

Clients that support only WPA2-PSK send an authentication frame with authentication algorithm set to open.

ArubaOS initiates a 4-way key handshake similar to WPA2.

WPA3-Enterprise

WPA3-Enterprise enforces top secret security standards for an enterprise Wi-Fi in comparison to secret security standards. Top secret security standards includes:

- Deriving at least 384-bit PMK/MSK using Suite B compatible EAP-TLS.
- Securing pairwise data between STA and authenticator using AES-GCM-256.
- Securing group addressed data between STA and authenticator using AES-GCM-256.
- Securing group addressed management frames using BIP-GMAC-256



ArubaOS supports WPA3-Enterprise only in non-termination 802.1X, tunnel-forward, and decrypt-tunnel modes. WPA3-Enterprise compatible 802.1X authentication occurs between STA and CPPM.

WPA3-Enterprise advertises or negotiates the following capabilities in beacons, probes response, or 802.11 association:

- AKM Suite Selector as 00-0F-AC:12
- Pairwise Cipher Suite Selector as 00-0F-AC:9
- Group data cipher suite selector as 00-0F-AC:9
- Group management cipher suite (MFP) selector as 00-0F-AC:12

If WPA3-Enterprise is enabled, STA is successfully associated only if it uses one of the four suite selectors for AKM selection, pairwise data protection, group data protection, and group management protection. If a STA mismatches any one of the four suite selectors, the STA association fails.

WPA3 Opmodes

ArubaOS supports the WPA3-AES-CCM-128, WPA3-AES-GCM-256, WPA3-CNSA, and WPA3-SAE-AES opmodes. These opmodes work in tunnel and decrypt-tunnel modes and opmode transition is not applicable to WPA3-AES-CCM-128 and WPA3-CNSA opmodes.

The WPA3 opmodes are effective only on the 300 Series, 310 Series, 320 Series, 330 Series, 340 Series, 360 Series, 370 Series, 510 Series, 530 Series, and 550 Series access points. Other access points will reject these opmodes.

If your network topology includes a Mobility Master that runs ArubaOS 8.4.0.0 and managed devices that run ArubaOS 8.3.0.0 or earlier version, the WPA3 opmodes will be converted to their corresponding WPA2 versions (for example: WPA3-SAE-AES opmode will be converted to WPA2-PSK-AES).

Before using the WPA3-SAE-AES opmode, disable opmode-transition and configure a WPA hexkey or WPA passphrase as a pre-shared key. Use the WPA3 with SAE and PSK mode for SAE mixed mode

operation during transition. The opmode-transition is not applicable to WPA3-AES-CCM-128 and WPA3-CNSA opmodes.

WPA2-PSK-AES virtual APs will be not be automatically upgraded to WPA3-SAE-AES virtual APs. Hence, WPA2-PSK-AES virtual APs will not automatically work in mixed mode. Configure a WPA3-SAE-AES virtual AP with opmode-transition for the virtual AP to operate in mixed mode.

Management Frame Protection in WPA3

Management frame protection is supported only in tunnel mode in WPA3 opmodes and it is enabled by default. The **mfp-capable** and **mfp-required** parameters do not take effect when any WPA3 opmode is enabled.

Configuring WPA3

To configure WPA3, configure the **opmode** and **opmode-transition** parameters under the **wlan ssid-profile** command.

The **opmode-transition** parameter is enabled by default and provides backward compatibility for authentication and WPA3-SAE-AES opmode. Use the **opmode-transition** parameter as a fallback option if a client faces connectivity issues on the enhanced open authentication or WPA3-SAE-AES transition mode virtual APs.

The following CLI commands configure WPA3 opmode:

```
(host) [mynode] #configure terminal
(host) [mynode] (config) #wlan ssid-profile wpa3_mode
(host) [mynode] (SSID Profile "wpa3_mode) #opmode wpa3-aes-ccm-128
(host) [mynode] (SSID Profile "wpa3_mode) #opmode wpa3-aes-gcm-256
(host) [mynode] (SSID Profile "wpa3_mode) #opmode wpa3-cnsa
(host) [mynode] (SSID Profile "wpa3_mode) #opmode wpa3-sae-aes
```



Using the WPA3-SAE-AES opmode requires a pre-shared key. Configure either a WPA hexkey or WPA passphrase as a pre-shared key.

The following procedure describes how to configure WPA3 opmode:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN > SSID**.
3. To create a new SSID profile, click **+** and enter a name for the new SSID profile in **Profile name**.
4. Configure your SSID settings. The configuration parameters are described in WLAN SSID Profiles.
5. In **Encryption**, select **wpa3-aes-ccm-128**, **wpa3-aes-gcm-256**, **wpa3-cnsa**, or **wpa3-sae-aes**.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands disable WPA3 opmode:

```
(host) [mynode] #configure terminal
(host) [mynode] (config) #wlan ssid-profile wpa3_mode
(host) [mynode] (SSID Profile "wpa3_mode") #no opmode
```

The following procedure describes how to disable WPA3 opmode:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN > SSID**.

3. To create a new SSID profile, click **+** and enter a name for the new SSID profile in **Profile name**.
4. Configure your SSID settings. The configuration parameters are described in WLAN SSID Profiles.
5. In **Encryption**, unselect **wpa3-aes-ccm-128**, **wpa3-aes-gcm-256**, **wpa3-cnsa**, or **wpa3-sae-aes**.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands enable opmode transition:

```
(host) [mynode] #configure terminal
(host) [mynode] (config) #wlan ssid-profile wpa3_opmode
(host) [mynode] (SSID Profile "wpa3_opmode") #opmode-transition
```

The following procedure describes how to enable opmode transition:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN > SSID**.
3. To create a new SSID profile, click **+** and enter a name for the new SSID profile in **Profile name**.
4. Configure your SSID settings. The configuration parameters are described in WLAN SSID Profiles.
5. Select **Opmode transition**.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands disable opmode transition:

```
(host) [mynode] #configure terminal
(host) [mynode] (config) #wlan ssid-profile wpa3_opmode
(host) [mynode] (SSID Profile "wpa3_opmode") #no opmode-transition
```

The following procedure describes how to disable opmode transition:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN > SSID**.
3. To create a new SSID profile, click **+** and enter a name for the new SSID profile in **Profile name**.
4. Configure your SSID settings. The configuration parameters are described in WLAN SSID Profiles.
5. Unselect **Opmode transition**.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

WPA3 in Decrypt-Tunnel Mode

ArubaOS supports WPA3-Enterprise and WPA3-Personal in decrypt-tunnel mode.



In 200 Series, 210 Series, 220 Series and 270 Series access points have limited support for WPA3 in decrypt-tunnel mode, and only support Enhanced Open, SAE, and wap3-aes-ccm-128.

The following CLI commands configure WPA3 in decrypt-tunnel mode:

```
(host) [mynode] #configure terminal
(host) [mynode] (config) #wlan virtual-ap wpa3_dtunnel_mode
(host) [mynode] (Virtual AP profile "wpa3_dtunnel_mode") #forward-mode decrypt-tunnel
```

```
(host) [mynode] (Virtual AP profile "wpa3_dtunnel_mode") #wlan ssid-profile wap3-dtunnel_
test
(host) [mynode] (SSID Profile "wpa3_dtunnel_test") #opmode wpa3-sae-aes
```

Fast BSS Transition Support for WPA3

ArubaOS supports Fast BSS Transition (802.11r) for the following WPA3 modes in both tunnel-forwarding and decrypt-tunnel modes for all APs which support WPA3:

- WPA3-Personal – SAE
- WPA3-Personal – SAE/WPA2-PSK Mixed mode
- WPA3-Enterprise Basic option
- WPA3-Enterprise non-CNSA mode with GCMP-256 Cipher Suite

Chapter 15
Stateful and WISPr Authentication

Mobility Master supports stateful 802.1X authentication, stateful NT LAN Manager authentication, and authentication for WISPr. Stateful authentication differs from 802.1X authentication in that Mobility Master does not manage the authentication process directly, but instead monitors the authentication messages between a user and an external authentication server, 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 on page 311
- WISPr Authentication on page 312
- Stateful Authentication Best Practices on page 312
- Configuring Stateful 802.1X Authentication on page 312
- Configuring Stateful NT LAN Manager Authentication on page 313
- Configuring Stateful Kerberos Authentication on page 315
- Configuring WISPr Authentication on page 316

Stateful Authentication

Mobility Master supports three different types of stateful authentication:

- **Stateful 802.1X authentication**: This feature allows Mobility Master to learn the identity and role of a user connected to a third-party AP, and is useful for authenticating users to networks with APs from multiple vendors. When an 802.1X-capable access point sends an authentication request to a RADIUS server, Mobility Master inspects this request and the associated response to learn the authentication state of the user. It then applies an identity-based user role through the Policy Enforcement Firewall.
- **Stateful Kerberos authentication**: Stateful Kerberos authentication configures Mobility Master to monitor the Kerberos authentication messages between a client and a Windows authentication server. If the client successfully authenticates through a Kerberos authentication server, Mobility Master recognizes that the client has been authenticated and assigns that client a specified user role.
- **Stateful NT LAN Manager authentication**: NT LAN Manager is a suite of Microsoft authentication and session security protocols. You can configure Mobility Master to monitor the NT LAN Manager authentication messages between a client and a Windows authentication server. If the client successfully authenticates through an NT LAN Manager authentication server, Mobility Master recognizes that the client has been authenticated and assigns that client a specified user role.

  The default Windows authentication method has changed from the older NT LAN Manager protocol to the newer Kerberos protocol, starting with Windows 2000. Therefore, stateful NT LAN Manager authentication is most useful for networks with legacy, pre-Windows 2000 clients. Also note that unlike other types of authentication, all users authenticated through stateful NT LAN Manager authentication must be assigned to the user role specified in the Stateful NT LAN Manager Authentication profile. The Aruba stateful NT LAN Manager authentication does not support placing users in various roles based upon group membership or other role-derivation attributes.

WISPr Authentication

WISPr authentication allows a smart client to authenticate to the network when roaming between wireless Internet service providers, even if the wireless hotspot uses an ISP, for which the client may not have an account.

If you are a hotspot operator using WISPr authentication, and a client that has an account with your ISP attempts to access the Internet at your hotspot, your ISP's WISPr AAA server authenticates that client directly and allows the client to access the network. If, however, the client only has an account with a partner ISP, your ISP's WISPr AAA server forwards that client's credentials to the partner ISP's WISPr AAA server for authentication. Once the client has been authenticated on the partner ISP, it is authenticated on your hotspot's own ISP, as per their service agreements. After your ISP sends an authentication message to Mobility Master, the default WISPr user role is assigned to that client.

Mobility Master supports the following smart clients, which enable client authentication and roaming between hotspots by embedding iPass Generic Interface Specification *redirect*, *proxy*, *authentication*, and *logoff* messages within HTML messages:

- iPass
- Boingo
- Trustive
- weRoam
- AT&T

Stateful Authentication Best Practices

Before you can configure a stateful authentication feature, you must define the user role you want to assign to the authenticated users and create a server group, which includes a RADIUS authentication server for stateful 802.1X authentication or a Windows server for stateful NT LAN Manager authentication. For details on performing these tasks, refer to the following sections of this User Guide:

- Roles and Policies on page 491
- Configuring Authentication Servers on page 193
- Configuring a Windows Server on page 206
- Configuring Server Groups on page 208

You can use the default stateful NT LAN Manager authentication and WISPr authentication profiles to manage the settings for these features, or you can create additional profiles as desired. 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 cannot configure more than one Stateful 802.1X profile.

Configuring Stateful 802.1X Authentication

When configuring 802.1X authentication for clients on non-Aruba APs, you must specify the group of RADIUS servers that performs user authentication and assign roles to users who successfully complete authentication. When the user logs off or shuts down the client machine, Mobility Master notes the deauthentication message from the RADIUS server and changes the user's role from the specified authenticated role back to the login role. For details on defining a RADIUS server used for stateful 802.1X authentication, see Configuring Authentication Servers on page 193.

The following procedure describes how to configure the Stateful 802.1X Authentication profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Authentication** page.
2. Under the **L2 Authentication** tab, select **Navigation** > **Stateful 802.1X Authentication**.
3. Select the role assigned to stateful 802.1X authenticated users from the **Default Role** drop-down list.
4. Specify the **Timeout** period for authentication requests, between 1 and 20 seconds.

   The default value is 10 seconds.
5. Select the **Mode** check box to enable stateful 802.1X authentication.
6. Click **Submit**.
7. Select **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure stateful 802.1X authentication. The first set of commands defines the RADIUS server used for 802.1X authentication, and the second set assigns that server to a server group. The third set associates the server group with the stateful 802.1X authentication profile, then sets the authentication role and timeout period.

```
(host) [md] (config) #aaa authentication-server radius <rad-server-name>
acctport <acctport>
authport <authport>
clone <source>
enable
enable-ipv6
enable-radsec
host <host>
key <key>
nas-identifier <nas-identifier>
nas-ip <nas-ip>
retransmit <retransmit>
timeout <timeout>
use-ip-for-calling-station
use-md5
```

```
(host) [md] (config) #aaa server-group <sg_name>
allow-fail-through
auth-server <name> [match-authstring {contains <sub_string>|equals <sub_
string>|starts-with <sub_string>][match-fqdn {all|<fqdn>}][position <prio>][trim-fqdn]
clone <source>
load-balance
set {role|vlan} condition <attribute> [contains <operand>|ends-with <operand>|equals
<operand>|not-equals <operand>|starts-with <operand>][value-of][set-value <set-value-
str>][position <number>]
```

```
(host) [md] (config) #aaa authentication stateful-dot1x
default-role <default-role>
enable
server-group <srv-group>
timeout <timeout>
```

The following CLI commands display the servers and profiles configured for stateful 802.1X authentication:

```
(host) [md] #show aaa authentication-server radius
(host) [md] #show aaa server-group
(host) [md] #show aaa authentication stateful-dot1x
```

Configuring Stateful NT LAN Manager Authentication

The Stateful NTLM Authentication profile requires that you specify a server group, which includes the servers performing NT LAN Manager authentication and the role to be assigned to users who are successfully authenticated. For details on defining a windows server used for NT LAN Manager authentication, see Configuring a Windows Server on page 206.

When a user logs off or shuts down the client machine, the user remains in the authenticated role until the user ages out, meaning there is no user traffic for the amount of time specified in the **User idle timeout** setting under **Configuration > Authentication > Advanced > Authentication Timers**.

The following procedure describes how to configure a stateful NTLM authentication profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication** page.
2. Select **Stateful NTLM Authentication** from the **L3 Authentication** tab.
3. Under **Stateful NTLM Authentication Profile: New Profile**, click the **+** to add a new profile entry. To modify an existing stateful NT LAN Manager authentication profile, select a profile entry below **Stateful NTLM Authentication** in the **L3 Authentication** list.
4. Enter a **Profile name**.
5. From the **Default Role** drop-down list, select the role to be assigned to all users after completing stateful NT LAN Manager authentication.
6. Select the **Mode** check box to enable stateful NT LAN Manager authentication.
7. Specify the **Timeout** period for authentication requests, between 1 and 20 seconds.

   The default value is 10 seconds.
8. Click **Submit**.
9. In the **L3 Authentication** list, select the **Server Group** entry below the stateful NT LAN Manager authentication profile.
10. Select the group of Windows servers to be used for stateful NT LAN Manager authentication from the **Server Group** drop-down list.
11. To enable authentication fail through and load balancing, select the check boxes for **Fail Through** and **Load Balance**.
12. Click **Submit**.
13. Select **Pending Changes**.
14. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure stateful NT LAN Manager authentication. The first set of commands defines the Windows server used for NT LAN Manager authentication, and the second set adds that server to a server group. The third set associates that server group with the stateful NT LAN Manager authentication profile, then defines the profile settings.

```
(host) [md] (config) #aaa authentication-server windows <windows_server_name>
clone <source>
domain <domain>
enable
host <host>
```

```
(host) [md] (config) #aaa server-group <sg_name>
allow-fail-through
auth-server <name> [match-authstring {contains <sub_string>|equals <sub_
string>|starts-with <sub_string>][match-fqdn {all|<fqdn>}][position <prio>][trim-fqdn]
clone <source>
load-balance
set {role|vlan} condition <attribute> [contains <operand>|ends-with <operand>|equals
<operand>|not-equals <operand>|starts-with <operand>][value-of][set-value <set-value-
str>][position <number>]
```

```
(host) [md] (config) #aaa authentication stateful-ntlm <profile-name>
clone <source>
default-role <default-role>
enable
server-group <server-group>
timeout <timeout>
```

The following CLI commands display the servers and profiles configured for stateful NT LAN Manager authentication:

```
(host) [md] #show aaa authentication-server window
(host) [md] #show aaa server-group
(host) [md] #show aaa authentication stateful-ntlm
```

Configuring Stateful Kerberos Authentication

The Stateful Kerberos Authentication profile requires that you specify a server group, which includes the Kerberos servers and the role assigned to authenticated users. For details on defining a windows server used for Kerberos authentication, see Configuring a Windows Server on page 206.

When the user logs off or shuts down the client machine, the user remains in the authenticated role until the user ages out, meaning there is no user traffic for the amount of time specified in the **User idle timeout** setting under **Configuration** > **Authentication** > **Advanced** > **Authentication Timers**

The following procedure describes how to configure a stateful Kerberos authentication profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Authentication** page.
2. Select **Stateful Kerberos Authentication** from the **L3 Authentication** tab.
3. Under **Stateful Kerberos Authentication Profile: New Profile**, click the **Add** button to add a new profile entry.

   To modify an existing stateful Kerberos authentication profile, select a profile entry below **Stateful Kerberos Authentication** in the **L3 Authentication** list.
4. Enter a **Profile name**.
5. From the **Default Role** drop-down list, select the role to be assigned to all users after completing stateful Kerberos authentication.
6. Specify the **Timeout** period for authentication requests, between 1 and 20 seconds.

   The default value is 10 seconds.
7. Click **Submit**.
8. In the **All Profiles** list, select the **Server Group** entry below the stateful Kerberos authentication profile.
9. Select the group of Windows servers to be used for stateful Kerberos authentication from the **Server Group** drop-down list.
10. To enable authentication fail through and load balancing, select the check boxes for **Fail Through** and **Load Balance**.
11. Click **Submit**.
12. Select **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure stateful Kerberos authentication. The first set of commands defines the server used for Kerberos authentication, and the second set adds that server to a server

group, and the third set of commands associates that server group with the stateful NT LAN Manager authentication profile then defines the profile settings.

```
(host) [md] (config) #aaa authentication-server windows <windows_server_name>
clone <source>
domain <domain>
enable
host <host>
```

```
(host) [md] (config) #aaa server-group <sg_name>
allow-fail-through
auth-server <name> [match-authstring {contains <sub_string>|equals <sub_
string>|starts-with <sub_string>][match-fqdn {all|<fqdn>}][position <prio>][trim-fqdn]
clone <source>
load-balance
set {role|vlan} condition <attribute> [contains <operand>|ends-with <operand>|equals
<operand>|not-equals <operand>|starts-with <operand>][value-of][set-value <set-value-
str>][position <number>]
```

```
(host) [md] (config) #aaa authentication stateful-kerberos <profile-name>
clone <source>
default-role <default-role>
server-group <server-group>
timeout <timeout>
```

The following CLI commands display the servers and profiles configured for stateful Kerberos authentication:

```
(host) [md] #show aaa authentication-server windows
(host) [md] #show aaa server-group
(host) [md] #show aaa authentication stateful-kerberos
```

Configuring WISPr Authentication

The WISPr authentication profile includes parameters to define RADIUS attributes, default roles for authenticated WISPr users, the maximum number of authentication failures, and login wait times. The WISPr-Location-ID, sent from Mobility Master to the WISPr RADIUS server, is the concatenation of the International Organization for Standardization Country Code, E.164 Country Code, E.164 Area Code, and SSID or Zone parameters configured in this profile.

The parameters used 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 International Organization for Standardization and International Telecommunication Union country and area codes at the International Organization for Standardization and International Telecommunication Union websites (www.iso.org) and http://www.itu.int.)

The following procedure describes how to configure a WISPr authentication profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication** page.
2. Select **WISPr Authentication** from the **L3 Authentication** tab.
3. Under **WISPr Authentication Profile: New Profile**, click the **+** to add a new profile entry.

   To modify an existing WISPr authentication profile, select a profile entry below **WISPr Authentication** in the **All Profiles** list.
4. Enter a **Profile name**.
5. Define values for the following parameters:

Table 64: *WISPr Authentication Profile Parameters*

Parameter	Description
Default Role	Default role assigned to users that complete WISPr authentication.
Max Authentication failures	Maximum number of failed WISPr authentication attempts permitted for each user.
User Agent String	User agent that identifies and provides details on the browser used during an HTTP request
Logon wait minimum wait	If the controller's CPU utilization has surpassed the **Login wait CPU utilization threshold value**, the **Logon wait minimum wait** parameter defines the minimum number of seconds a user has to wait to retry a login attempt. Range: 1–10 seconds. Default: 5 seconds.
Logon wait maximum wait	If the controller's CPU utilization has surpassed the **Login wait CPU utilization threshold** value, the **Logon wait maximum wait** parameter defines the maximum number of seconds a user has to wait to retry a login attempt. Range: 1–10 seconds. Default: 10 seconds.
Logon wait CPU utilization threshold	Percentage of CPU utilization at which the maximum and minimum login wait times are enforced. Range: 1–100%. Default: 60%.
WISPr Location-ID ISO Country Code	The International Organization for Standardization 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 or Zone section of the WISPr Location ID.
WISPr Operator Name	Name identifying the hotspot operator.
WISPr Location Name	Name identifying the hotspot location. If no name is defined, the parameter uses the name of the associated AP.

6. Click **Submit**.
7. In the **All Profiles** list, select the **Server Group** entry below the WISPR authentication profile.
8. Select the group of RADIUS servers to be used for WISPr authentication from the **Server Group** drop-down list.
9. To enable authentication fail through and load balancing, select the check boxes for **Fail Through** and **Load Balance**.
10. Click **Submit**.
11. Select **Pending Changes**.
12. In the **Pending Changes** window, select the check box and click **Deploy changes**.



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.

The following CLI commands configure WISPr authentication. The first set of commands defines the RADIUS server used for WISPr authentication, and the second set adds that server to a server group. The third set of commands associates that server group with the WISPr authentication profile, then defines the profile settings.

```
(host) [md] (config) #aaa authentication-server radius <rad-server-name>
acctport <acctport>
authport <authport>
clone <source>
enable
enable-ipv6
enable-radsec
host <host>
key <key>
nas-identifier <nas-identifier>
nas-ip <nas-ip>
retransmit <retransmit>
timeout <timeout>
use-ip-for-calling-station
use-md5
```

```
(host) [md] (config) #aaa server-group <sg_name>
allow-fail-through
auth-server <name> [match-authstring {contains <sub_string>|equals <sub_
string>|starts-with <sub_string>][match-fqdn {all|<fqdn>}][position <prio>][trim-fqdn]
clone <source>
load-balance
set {role|vlan} condition <attribute> [contains <operand>|ends-with <operand>|equals
<operand>|not-equals <operand>|starts-with <operand>][value-of][set-value <set-value-
str>][position <number>]
```

```
(host) [md] (config) #aaa authentication wispr <profile-name>
agent_string <agent_string>
clone <source>
default-role <default-role>
logon-wait {cpu-threshold <cpu-threshold>|maximum-delay <maximum-delay>|minimum-delay
<minimum-delay>}
max-authentication-failures <max-authentication-failures>
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-operator-name>
```

The following CLI commands display the servers and profiles configured for WISPr authentication:

```
(host) [md] #show aaa authentication-server radius
(host) [md] #show aaa server-group
(host) [md] #show aaa authentication wispr
```

Chapter 16
Certificate Revocation

The Certificate Revocation feature enables the Mobility Master or the Managed Device to perform real-time certificate revocation checks using the OCSP, or traditional certificate validation using the CRL client.

Topics in this chapter include:

- Understanding OCSP and CRL on page 319
- Configuring the Mobility Master or Managed Device as an OCSP Client on page 320
- Configuring the Mobility Master or Managed Device as a CRL Client on page 321
- Configuring the Mobility Master or Managed Device as an OCSP Responder on page 323
- Certificate Revocation Checking for SSH Pubkey Authentication on page 324

Understanding OCSP and CRL

OCSP (RFC 2560) is a standard protocol that consists of an OCSP client and an OCSP responder. This protocol determines revocation status of a given digital public-key certificate without downloading the entire CRL.

CRL is the traditional method of checking certificate validity. A CRL provides a list of certificate serial numbers that have been revoked or are no longer valid. CRLs let the verifier check the revocation status of the presented certificate while verifying it. CRLs are limited to 512 entries.

Both the Delegated Trust Model and the Direct Trust Model are supported to verify digitally signed OCSP responses. Unlike the Direct Trust Model, the Delegated Trust Model does not require the OCSP responder certificates to be explicitly available on the Mobility Master or the managed device.

This section describes the following topics:

Configuring the Mobility Master or the Managed Device as OCSP and CRL Clients

The Mobility Master or the managed device can act as an OCSP client and issue OCSP queries to remote OCSP responders located on the intranet or Internet. Since many applications in ArubaOS (such as IKE), use digital certificates, a protocol such as OCSP needs to be implemented for revocation.

An entity that relies on the content of a certificate (a relying party) needs to check before accepting the certificate as valid. Once it is verified that the certificate has not been revoked, the OCSP client retrieves certificate revocation status from an OCSP responder. The responder may be the CA that has issued the certificate in question, or it may be some other designated entity which provides the service on behalf of the CA. A *revocation checkpoint* is a logical profile that is tied to each CA certificate that the Mobility Master or the managed device has (trusted or intermediate). Also, the user can specify revocation preferences within each profile.

The OCSP request is not signed by the Aruba OCSP client at this time. However, the OCSP response is always signed by the responder.

Both OCSP and CRL configuration and administration is usually performed by the administrator who manages the web access policy for an organization.

In small networks where there are is no Internet connection or connection to an OCSP responder, CRL is preferable to than OCSP.

Configuring the Mobility Master or Managed Device as an OCSP Responder

The Mobility Master or the managed device can be configured to act as an OCSP responder (server) and respond to OCSP queries from clients that want to obtain revocation status of certificates.

The OCSP responder on the Mobility Master or the managed device is accessible over HTTP port 8084. You cannot configure this port. Although the OCSP responder accepts signed OCSP requests, it does not attempt to verify the signature before processing the request. Therefore, even unsigned OCSP requests are supported.

The Mobility Master or the managed device as an OCSP responder provides revocation status information to Aruba applications that use CRLs. This is useful in small disconnected networks where clients cannot reach outside OCSP server to validate certificates. Typical scenarios include client to client or client to other server communication situations where the certificates of either party need to be validated.

Configuring the Mobility Master or Managed Device as an OCSP Client

When OCSP is used as the revocation method, you need to configure the OCSP responder certificate and the OCSP URL.

You can configure the Mobility Master or Managed Device as an OCSP client using the WebUI.

In the WebUI

Perform the following steps to configure Mobility Master or Managed Device as an OCSP client:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Certificates** tab.
2. Expand the **Import Certificates** accordion.
3. Click **+** in the **Import Certificates** section.
4. Enter the following certificate details in the **New Certificate** section:
    a. Enter a name in the **Certificate name** text box. This name identifies the certificate you are importing.
    b. Enter the certificate filename in the **Certificate filename** text box. Click the **Browse** button to enter the full pathname.
    c. Enter a password in the **Optional passphrase** text box. The password is optional.
    d. If you opted to use the optional password (in step c), re-enter the password in the **Retype passphrase** text box.
    e. Select a certificate format from the **Certificate format** drop-down list. You can import certificates of format DER, P12, PEM, PFX, PKCS12, and PKCS7.

f. Select **OCSPResponderCert** from the **Certificate type** drop-down list.



A revocation check method (OCSP or CRL) can be chosen independently for every revocation checkpoint. In this example, we are only describing the OCSP check method.

When this certificate is imported, it is maintained in the certificate store for OCSP responder certificates. These certificates are used for signature verification.

5. Click **Submit**. The certificate appears in the **Import Certificates** section.
6. For detailed information about an imported certificate, click the certificate from the certificate list.
7. Click the **Revocation Checkpoint** accordion menu.
8. In the **Revocation Checkpoint** section, click the record for which you want to configure the revocation checkpoint. The **Revocation Checkpoint > <RCP name>** section is displayed.
   a. Select **ocsp** from the **Revocation method 1** drop-down list as the primary check method.
   b. In the **OCSP URL** text box, enter the URL of the OCSP responder.
   c. Select the OCSP certificate that you want to configure from the **OCSP responder cert** drop-down list.
9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box indicating the pending change and click **Deploy Changes**.

    You can configures an OCSP client with the revocation check method as OCSP for a revocation checkpoint using the CLI.

In the CLI

The OCSP responder certificate is configured first . The corresponding OCSP responder service is available at http://10.4.46.202/ocsp. The check method is OCSP for the revocation checkpoint.

Execute the following command to configures an OCSP client with the revocation check method as OCSP for a revocation checkpoint.

```
(host)[mynode](config) #crypto-local pki rcp <name>
(host)[mynode](config-submode) #ocsp-responder-cert <ocsp_responder_cert>>
(host)[mynode](config-submode) #ocsp-url http://10.4.46.202/ocsp
(host)[mynode](config-submode) #revocation-check ocsp
```

The **show crypto-local pki OCSPResponderCert** command lists the contents of the OCSP Responder Certificate store.

The **show crypto-local pki rcp <rcp_name>** command shows the entire configuration for a given revocation checkpoint.

Configuring the Mobility Master or Managed Device as a CRL Client

CRL is the traditional method of checking certificate validity. When you want to check certificate validity using a CRL, import the CRL. You can import CRLs only by using the WebUI.

In the WebUI

Perform the following steps to configure the Mobility Master as a CRL client:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Certificates** tab.
2. Expand the **Import Certificates** accordion.
3. Click **+** in the **Import Certificates** section.
4. Enter the following certificate details in the **New Certificate** section:
   a. Enter a name in the **Certificate name** text box. This name identifies the certificate you are importing.
   b. Enter the certificate filename in the **Certificate filename** text box. Click the **Browse** button to enter the full pathname.
   c. Enter a password in the **Optional passphrase** text box. The password is optional.
   d. If you opted for using the optional password (in step c), re-enter the password in the **Retype passphrase** text box.
   e. Select a certificate format from the **Certificate format** drop-down list. You can import certificates of format DER, P12, PEM, PFX, PKCS12, and PKCS7.
   f. Select **CRL** from the **Certificate type** drop-down list.

   A revocation check method (OCSP or CRL) can be chosen independently for every revocation checkpoint. In this example, we are only describing the CRL check method.

   When this CRL is imported, it is maintained in the store for CRLs. These CRLs are used for signature verification.
5. Click **Submit**. The CRL appears in the **Import Certificates** section.
6. For detailed information about an imported CRL, click the CRL from the CRL list.
7. Click the **Revocation Checkpoint** accordion menu.
   a. In the **Revocation Checkpoint** section, click the record for which you want to configure the revocation checkpoint. The **Revocation Checkpoint > <RCP name>** section is displayed.
   b. Select **crl** from the **Revocation method 1** drop-down list.
   c. In the **CRL location** text box, enter the CRL you want to use for this revocation checkpoint. The CRLs listed are files that have already been imported onto the Mobility Master or the managed device.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box indicating the pending change and click **Deploy Changes**.

    You can configure an OCSP responder with the check method as CRL for a revocation check point using the CLI.

In the CLI

Run the following command to configures an OCSP responder with the check method as CRL for a revocation check point:

```
(host)[mynode](config) #crypto-local pki rcp <rcp-name>
(host)[mynode](config-submode) #crl-location file <filename>
(host)[mynode](config-submode) #revocation-check crl
```

Configuring the Mobility Master or Managed Device as an OCSP Responder

When configured as an OCSP responder, the Mobility Master or the managed device provides revocation status information to ArubaOS applications that use CRLs.

You can configure Mobility Master or Managed Device as an OCSP responder using the WebUI or the CLI.

In the WebUI

Perform the following steps to configure the Mobility Master as an OCSP responder:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Certificates** tab.
2. Expand the **Import Certificates** accordion.
3. Click **+** in the **Import Certificates** section.
4. Enter the following certificate details in the **New Certificate** section:
   a. Enter a name in the **Certificate name** text box. This name identifies the certificate you are importing.
   b. Enter the certificate filename in the **Certificate filename** text box. Click the **Browse** button to enter the full pathname.
   c. Enter a password in the **Optional passphrase** text box. The password is optional.
   d. If you opted for using the optional password (in step c), re-enter the password in the **Retype passphrase** text box.
   e. Select a certificate format from the **Certificate format** drop-down list. You can import certificates of format DER, P12, PEM, PFX, PKCS12, and PKCS7.
   f. Select **OCSPSignerCert** from the **Certificate type** drop-down list.

   - When this certificate is imported, it is maintained in the certificate store for OCSP signer certificates. These certificates are used for signature verification.
   - The OCSP signer cert signs OCSP responses for this revocation checkpoint. The OCSP signer cert can be the same trusted CA as the checkpoint, a designated OCSP signer certificate issued by the same CA as the checkpoint or some other local trusted authority.
   - If you do not specify an OCSP signer cert, OCSP responses are signed using the global OCSP signer certificate. If that is not present, than an error message is sent out to clients.



The OCSP signer certificate takes precedence over the global OCSP signer certificate as it is checkpoint specific.

5. Click **Submit**. The certificate appears in the **Import Certificates** section.
6. For detailed information about an imported certificate, click the certificate from the certificate list.
7. Click the **Revocation Checkpoint** accordion menu.
   a. Click the **Enable OCSP responder** toggle switch to enable this setting.
   b. **Enable OCSP responder** is a global option that turns the OCSP responder service on or off on the Mobility Master or the managed device. The default is disabled (off). Enabling this option automatically adds the OCSP responder port (TCP 8084) to the permit list in the CP firewall so this can be accessed from outside the Mobility Master or the managed device.

c. Select the **OCSPSignerCert** to be used to sign OCSP responses for this revocation checkpoint from the **OCSP certificates** drop-down list .

d. In the **Revocation Checkpoint** section, click the record for which you want to configure the revocation checkpoint. The **Revocation Checkpoint > <RCP name>** section is displayed.

e. Select **ocsp** from the **Revocation method 1** drop-down list as the primary check method. Optionally, select a backup check method from the **Revocation method 2** drop-down list.

f. In the **CRL location** text box, enter the CRL you want used for this revocation checkpoint. The CRLs listed are files that have already been imported onto the Mobility Master or the managed device.

g. Click the **Enable OCSP responder** toggle switch to enable this setting.

h. Select **OCSPSignerCert** from the **OCSP signer cert** drop-down list.

8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box indicating the pending change and click **Deploy Changes**.

In the CLI

Run the following commands to configure the Mobility Master or the managed device as an OCSP responder.

```
(host)[mynode](config) #crypto-local pki service-ocsp-responder
(host)[mynode](config) #crypto-local pki rcp <name>
(host)[mynode](config-submode) #ocsp-signer-cert oscsp_CA1
(host)[mynode](config-submode) #crl-location file <filename>
(host)[mynode](config-submode) #enable-ocsp-responder
```

Certificate Revocation Checking for SSH Pubkey Authentication

This feature allows the ssh-pubkey management user to be optionally configured with a revocation checkpoint. This meets the requirement for a two-factor authentication and integration of device management with PKI for SSH pubkey authentication. The ArubaOS implementation of SSH using Pubkey authentication is designed for integration with smart cards or other technologies that use X.509 certificates.

The revocation checkpoint checks the revocation status of the SSH user's client certificate before permitting access. If the revocation check fails, the user is denied access using the ssh-pubkey authentication method. However, the user can still authenticate through a username and password if configured to do so.

For information about configuring a revocation checkpoint, see Certificate Revocation on page 319.

Configuring the SSH Pubkey User with Revocation Checkpoint

You can configure the SSH pubkey user with revocation checkpoint to check the validity of the user's X.509 certificate, using the WebUI or the CLI.

In the WebUI

Perform the following steps to configure the SSH Pubkey User with revocation checkpoint:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Admin** tab.
2. Expand the **Management User** accordion.
3. Click the **Show users with certificate authentication** link. The **Management Users with Certificate Authentication** section is displayed.
4. In the **Management Users with Certificate Authentication** section, click **+**. The **Management Users with Certificate Authentication > New User** section is displayed.
5. In the **Management Users with Certificate Authentication > New User** section, perform the following steps:
   a. Select **WebUI & CLI through SSH** from the **Interface to connect** drop-down list.
   b. Enter a username in the **Username** text box and select a node from the **Node** drop-down list.
   c. Select a role from the **Role** drop-down list.
   d. Select a certificate from the **Trusted CA certificate name** drop-down list.
   e. Enter the client certificate serial number in the **Client certificate serial no.** text box.
   f. For **Authentication server**, select either **Internal server or External Server**. Select **External server** if you want the user to be authenticated by an external authentication server. When this option is selected, **Role** and **Client certificate serial no fields** disappear.
   g. Select a **Role** and **Client Certificate** in the **CLI through SSH section**.
6. To specify the revocation checkpoint, perform either of the following tasks:
7. To enable revocation checkpoint, select a valid configured revocation checkpoint from **Revocation checkpoint** drop-down list.
8. Select **None** if you do not want the revocation checkpoint enabled for the SSH pubkey user.
9. Click **Submit**.
10. Click **Pending Changes** at the top of the window.
11. In the **Pending Changes** window, select the check box indicating the pending change and click **Deploy Changes**.

In the CLI

The CLI allows you to configure an optional revocation checkpoint for an ssh-pubkey user. Users can still be configured without the revocation checkpoint. In this example, the certificate name is “client1-rg,”, the username is “test1,” the role name is “root,” and the revocation checkpoint is “ca-rg:”

```
(host)[mynode](config) #mgmt-user ssh-pubkey client-cert  client1-rg test1 root rcp
ca-rgg
```

In this example, a user is configured without the revocation checkpoint:

```
(host)[mynode](config) #mgmt-user ssh-pubkey client-cert client2-rg test2 root
```

Displaying Revocation Checkpoint for the SSH Pubkey User

The revocation checkpoint checks the revocation status of the SSH user’s client certificate before permitting access. If the revocation check fails, the user is denied access using the ssh-pubkey authentication method. However, the user can still authenticate through a username and password if configured to do so. This feature allows the ssh-pubkey management user to be optionally configured with a revocation checkpoint. This meets the requirement for a two-factor authentication and integration of device management with PKI for SSH pubkey authentication. The ArubaOS

implementation of SSH using Pubkey authentication is designed for integration with smart cards or other technologies that use X.509.

The column **REVOCATION CHECKPOINT** displays the configured revocation checkpoint for the ssh-pubkey user. If no revocation checkpoint is configured for the user, the entry **none** is displayed.

You can view the revocation checkpoint using the WebUI or the CLI.

In the WebUI

Perform the following steps to view the revocation checkpoint for an SSH pubkey user:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Admin**.
2. Expand the **Management User** accordion.
3. Click the **Show users with certificate authentication** link. The **Management Users with Certificate Authentication** section is displayed.
4. Click the user name for which you want to know the configured revocation checkpoint. The **Management User > <username>** section is displayed.
5. The **Revocation Checkpoint** column displays the revocation checkpoint configured (if any) for the SSH pubkey user.

In the CLI

Run the following command to view the revocation checkpoint from the Mobility Master node hierarchy:

```
(host)[mynode] #show mgmt-user ssh-pubkey
```

Removing the SSH Pubkey User

You can remove the SSH Pubkey user by using the WebUI or the CLI.

In the WebUI

Perform the following steps to remove the SSH Pubkey user:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Admin** tab.
2. Click the **Management User** accordion.
3. Click the **Show users with certificate authentication** link. The **Management Users with Certificate Authentication** section is displayed.
4. Click the username which you want to delete.
5. Click the bin box icon beside the management user that you want to delete.

In the CLI

Run the following command to remove the SSH Pubkey user from the Mobility Master node hierarchy:

```
(host)[mynode](config) #no mgmt-user ssh-pubkey client-cert <certname> <username>
```

Chapter 17
Captive Portal Authentication

Captive portal is one of the methods of authentication supported by ArubaOS. A captive portal presents a web page which requires user action before network access is granted. The required action can be simply viewing and agreeing to an Acceptable Usage 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 Aruba VPN dialer for Microsoft VPN clients if the VPN is to be terminated on the Mobility Master.

The following list displays the key topics discussed:

- Captive Portal Authentication on page 327
- Understanding Captive Portal on page 327
- Configuring Captive Portal in the Base Operating System on page 329
- Configuring Captive Portal with a PEFNG License on page 331
- Sample Authentication with Captive Portal on page 334
- Configuring Guest VLANs on page 343
- Configuring Captive Portal Authentication Profiles on page 344
- Enabling Optional Captive Portal Configuration on page 349
- Personalizing the Captive Portal Page on page 353
- Creating Walled Garden Access on page 366
- Enabling Captive Portal Enhancements on page 368

Captive Portal supports the following deployment models:

- Mobility Master-Managed Device
- Stand-alone Controller

Mobility Master-Managed Device

Mobility Master is the root of a network hierarchy. A single Mobility Master oversees a number of managed devices that can be co-located or off-campus. In Mobility Master-Managed Device deployment model, all Captive Portal configuration is allowed only on the Mobility Master.

Stand-alone Controller

Captive Portal is supported in the stand-alone controller mode where the configuration can be performed on the controller irrespective of local or master controller.

Related Topics

Virtual Private Networks

Understanding Captive Portal

Understanding Captive Portal

You can configure captive portal for the following users:

- Guest users, where no authentication is required.
- Registered users, who must be authenticated against an external server or the internal database of the managed device.



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 also load up to 16 different customized login pages into the managed device. The login page displayed is based on the SSID to which the client associates.

Captive portal provides secure services to its users by using the following:

- Policy Enforcement Firewall Next Generation License
- Server Certificate

Policy Enforcement Firewall Next Generation License

The Policy Enforcement Firewall Next Generation License (PEFNG) license provides identity-based security for wired and wireless users through user roles and firewall rules. You can use captive portal with or without the PEFNG license installed in the Mobility Master. There are differences in how captive portal functions work and how you configure captive portal, depending on whether the license is installed.

Server Certificate

The Aruba managed device is designed to provide secure services through the use of digital certificates. The server certificate is installed on the managed device through the Mobility Master. A server certificate installed in the managed device verifies the authenticity of the managed devices for captive portal.

Aruba managed device ship with a demonstration self-signed certificate. Until you install a customer-specific server certificate in the managed device, this demonstration self-signed certificate is used by default for all secure HTTP connections such as captive portal. This self-signed 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 CA. You can generate a CSR on the managed device to submit to a CA.

The managed device can accept wild card server certificates (CN begins with an asterisk). If a wildcard certificate is uploaded (for example, CN=*.domain.com), the asterisk in CN is replaced with 'captiveportal-login' in order to derive the Captive Portal logon page URL (captiveportal-login.domain.com).

Once you have imported a server certificate from the Mobility Master to managed device, you can select the certificate to be used with captive portal.

Configuring Server Certificate

The following procedure describes how to select a certificate for captive portal:

1. Login to the Mobility Master.
2. In the **Managed Network**node hierarchy, navigate to the **Configuration > System > More > General** accordion.
3. Under **Captive Portal Certificate**, select the name of the imported certificate from the drop-down list.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands are used to select a certificate for captive portal:

```
(host) [mynode] #cd /md /<MAC_address>
(host) [<MAC_address>] (config) #web-server profile
(host) [<MAC_address>] (Web Server Configuration) #captive-portal-cert <certificate>
```

To specify a different server certificate for captive portal with the CLI, use the **no** command to revert to the default certificate *before* you specify the new certificate:

```
(host) [<MAC_address>] (config) #web-server profile
(host) [<MAC_address>] (Web Server Configuration) #captive-portal-cert ServerCert1
(host) [<MAC_address>] (Web Server Configuration) #no captive-portal-cert
(host) [<MAC_address>] (Web Server Configuration) #captive-portal-cert ServerCert2
```

Related Topics

Configuring Captive Portal in the Base Operating System

Configuring Captive Portal with a PEFNG License

Managing Certificates

Configuring Captive Portal in the Base Operating System

The base operating system (ArubaOS without any licenses) allows full network access to all users who connect to an ESSID, both guest and registered users. In the base operating system, you cannot configure or customize user roles; this function is only available by installing the PEFNG license. Captive portal allows you to control or identify who has access to network resources.

When you create a captive portal profile in the base operating system, an implicit user role is automatically created in the stand-alone controller 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.

In a Mobility Master-managed device topology, Mobility Master does not have the configuration which are related to PEFNG license, therefore the role is not created on the Mobility Master.



The WLAN Wizard within the ArubaOS WebUI allows for basic captive portal configuration for WLANs associated with the “default” ap-group: **Configuration > WLAN Wizard**. Follow the steps in the workflow pane within the wizard and refer to the help tab for assistance.

Following are the tasks for configuring captive portal in the base ArubaOS:

1. Create the Server Group name. In this example, the server group name is **cp-srv**.

   If you are configuring captive portal for registered users, configure the server(s) and create the server group. For more information about configuring authentication servers and server groups, see Authentication Servers on page 192.

2. 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 **c-portal** profile 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.

3. Create a 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. The initial role in the profile **aaa_c-portal** must be set to **c-portal**.

4. 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 created.

5. 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 CLI. Configuring the VLAN and authentication servers and server groups are described elsewhere in this document.

The following procedure describes how to configure captive portal in the base operating system:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > L3 Authentication** tab. Select **Captive Portal Authentication**.
   a. Click **+** in **Captive Portal Authentication Profile: New Profile**, enter a **Profile Name** (for example, **c-portal**).
   b. You can enable user login and guest login, and configure other captive portal profile parameters as described in Configuring Captive Portal Authentication Profiles on page 344.
   c. Click **Submit**.
3. 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 list.
   b. Click **Submit.**
4. Select the **AAA Profiles** tab.
   a. Expand **AAA Profiles**, click **+** in **AAA Profile: New Profile** to add a new profile. Enter a **Profile Name**(for example, **aaa_c-portal**), then click **Submit**.
   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 for stand-alone controller.



The Initial Role must be exactly the same as the name of the captive portal authentication profile you created.

d. Click **Submit**.

5. Navigate to the **Configuration > System > Profiles tab** and under Profiles, select **Wireless LAN**, then select **Virtual AP**.
6. To create a new virtual AP profile, Click + in **Virtual AP profile: New Profile**.
7. Enter the name for the virtual AP profile (for example, **vp_c-portal**). Make sure **Virtual AP enable** is selected.
8. For **VLAN**, enter the ID of the VLAN in which captive portal users are placed (for example, VLAN **20)**. Click **Submit**.
   a. In the **Profile Details** entry for the new virtual AP profile (**guestnet**), select the AAA profile you previously configured from the **AAA Profile** drop-down list and click **Submit**.
   b. In the **Profile Details** entry for the new virtual AP profile (**guestnet**), select the **SSID profile** and select a SSID profile from the **SSID** profile drop-down list.
   c. Enter the name for the **ESSID** profile (for example,**essid_c-portal**).
   d. For **Encryption**, select **opensystem**.
   e. At the bottom of the Profile Details page, click **Submit.**
9. Navigate to the **Configuration > AP Groups** page.
10. Select an AP Group and Click **WLANs** tab in the AP group window.
11. Click **+** under the WLANs tab and select the newly create virtual AP profile (guestnet) from the **Virtual-ap** drop-down list.
12. Click **Submit.**
13. Click **Pending Changes**.
14. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure captive portal in the base operating system:

```
(host) [md] (config) #aaa authentication captive-portal c-portal
server-group cp-srv
(host) [md] (config) #aaa profile aaa_c-portal
initial-role c-portal
(host) [md] (config) #wlan ssid-profile ssid_c-portal
essid c-portal-ap
(host) [md] (config) #wlan virtual-ap vp_c-portal
aaa-profile aaa_c-portal
ssid-profile ssid_c-portal
vlan 20
```

Related Topics

Configuring Captive Portal with a PEFNG License

Sample Authentication with Captive Portal

Configuring Captive Portal Authentication Profiles

Configuring Captive Portal with a PEFNG License

You must purchase and install the PEFNG license on the Mobility Master to use identity-based security features. There are two user roles that are important for captive portal:

- Default user role, which you specify in the captive portal authentication profile, is the role granted to clients upon captive portal authentication. This can be the predefined **guest** system role.
- Initial user role, which you specify in the AAA profile, directs clients who associate to the SSID to captive portal whenever the user initiates a Web browser connection. This can be the predefined **logon** system role.

  The captive portal authentication profile specifies the captive portal login page and other configurable parameters. The initial user role configuration must include the applicable captive portal authentication profile instance.



MAC-based authentication, if enabled on the Mobility Master, takes precedence over captive portal authentication.

Following are the basic tasks for configuring captive portal using role-based access provided by the Policy Enforcement Firewall software module:

1. Install the PEFNG license on the primary Mobility Master.

   For more information, see *Aruba Mobility Master Licensing Guide*.

2. Configure the user role for a default user.

   Create and configure user roles and policies for guest or registered captive portal users. For more information, see Configuring Policies and Roles on page 336.

3. Create a server group.

   If you are configuring captive portal for registered users, configure the server(s) and create the server group. For more information, see Authentication Servers on page 192

   If you are using the internal database of the managed device for user authentication, use the predefined “Internal” server group. The "internal" server is the local database on the Mobility Master. You need to configure entries in the internal database, as described in Authentication Servers on page 192.

4. 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. For more information, see Configuring Captive Portal Authentication Profiles on page 344.

5. Configure the initial user role.

   Create and configure the initial user role for captive portal. 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. For more information, see Modifying the Initial User Role on page 345.

6. Create the AAA Profile.

   Create and configure an instance of the AAA profile. Specify the initial user role. For more information, see Configuring the AAA Profile on page 345.

7. Create the SSID Profile. In this example, the profile name is **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.

8. Create the 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 WebUI and 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.

The following procedure describes how to configure captive portal with a PEFNG license:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > L3 Authentication** tab. Select the **Captive Portal Authentication** Profile.
   a. In the **Captive Portal Authentication Profile: New Profile** window, click **+** to create a new Captive Portal Authentication Profile and enter a **Profile Name** (for example, **c-portal**).
   b. Select the **Default role** (for example, **employee**) for captive portal users.
   c. Enable **Guest login** or **User login**, as well as other parameters (refer to *Captive Portal Authentication Profile Parameters* table).
   d. Click **Submit**.
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 list.
   b. Click **Submit.**
4. Select the **AAA Profiles** tab.
   a. Expand AAA Profiles and click **+** in the **AAA profile: New Profile** window to add a new profile. Enter a **Profile Name**(for example, **aaa_c-portal**).
   b. Set the **Initial role** to a role that you will configure with the captive portal authentication profile.
   c. Click **Submit**.
5. Navigate to the **Configuration > Roles and Policies> Roles** tab**.** Select a role and click + to add a new rule.
   a. To edit the predefined logon role, select the role and click + in the policies page that opens and select **Access Control**.
   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. Select the profile from the Captive Portal Profile drop-down list in Authentication tab under the selected Role.
   d. Click **Submit**.
6. Navigate to the **Configuration > AP Groups** page to configure the virtual AP profile.
7. Select the AP Group . Click **+** 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 list to create a new virtual AP profile. Enter the name for the virtual AP profile (for example, **vp_c-portal**), then click **Save.**
    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 **Save**.
    b. From the SSID profile drop-down list, 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. At the bottom of the Profile Details page, click **Submit.**
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, **900**).
    c. Click **Submit**.
11. Click **Pending Changes**.
12. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure captive portal with the PEFNG license:

```
(host) [md] (config) #aaa authentication captive-portal c-portal
default-role employee
server-group cp-srv
(host) [md] (config) #user-role logon
(host) [md] (config-submode)#access-list session c-portal
captive-portal c-portal
(host) [md] (config) #aaa profile aaa_c-portal
initial-role logon
(host) [md] (config) #wlan ssid-profile ssid_c-portal
essid c-portal-ap
vlan 900
(host) [md] (config) #wlan virtual-ap vp_c-portal
aaa-profile aaa_c-portal
ssid-profile ssid_c-portal
```

Related Topics

Configuring Captive Portal in the Base Operating System

Sample Authentication with Captive Portal

Configuring Captive Portal Authentication Profiles

Sample Authentication with Captive Portal

In the following example:

- Guest clients associate to the **guestnet** SSID which is an open WLAN. Guest clients are placed into VLAN 900 and assigned IP addresses by the internal DHCP server of the managed device. 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 internal database of the managed device. 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 with source network address translation.



This example assumes a Policy Enforcement Firewall Next Generation license is installed in the Mobility Master.

In this example, you create two user roles:

- **guest-logon** is a user role assigned to any client who associates to the guestnet SSID. Normally, any client that associates to an SSID will be placed into the *logon* system role. The **guest-logon** user role is more restrictive than the logon role.
- **auth-guest** is a user role granted to clients who successfully authenticate via the captive portal.

Creating a Guest User Role

The **guest-logon** user role consists of the following ordered policies:

- **captiveportal** is a predefined policy that allows captive portal authentication.
- **guest-logon-access** is a policy that you create with the following rules:
  - Allows DHCP exchanges between the user and the DHCP server during business hours while blocking other users from responding to DHCP requests.
  - Allows ICMP exchanges between the user and the managed devices during business hours.
- **block-internal-access** is a policy that you create that denies user access to the internal networks.



The **guest-logon** user role configuration needs to include the name of the captive portal authentication profile instance. You can modify the user role configuration after you create the captive portal authentication profile instance.

Creating an Auth-guest User Role

The **auth-guest** user role consists of the following ordered policies:

- **cplogout** is a predefined policy that allows captive portal logout.
- **guest-logon-access** is a policy that you create with the following rules:
  - Allows DHCP exchanges between the user and the DHCP server during business hours while blocking other users from responding to DHCP requests.
  - Allows DNS exchanges between the user and the public DNS server during business hours. Source-NAT the traffic using the IP interface of the managed devices 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. Source-NAT the traffic using the IP interface of the managed devices for the VLAN.
- Allows HTTP or HTTPS traffic from the user during business hours. Source-NAT the traffic using the I interface of the managed devices for the VLAN.

- **drop-and-log** is a policy that you create that denies all traffic and logs the attempted network access.

Configuring Policies and Roles

The following section describes how to configure roles and policies by creating a time range, creating aliases, creating a Guest-Logon-Access policy, Auth-Guest-Access policy, Block-Internal-Access policy, Drop-and-Log policy, and Auth-Guest role.

Creating a Time Range

The following procedure describes how to create the guest-logon-access policy:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Policies** tab.
3. Select **+** to add the guest-logon-access policy.
4. For **Policy Name**, enter **guest-logon-access.**
5. For **Policy Type**, select **Session**.
6. Click **Submit**.
7. Select the newly created **guest-logon-access** policy.
8. Click **+** under the **Policies > guest-logon-access** table.
9. In the **New Rule for guest-logon-access** popup, select **Access Control** option and click OK.
10. Under1 **guest-logon-access > New forwarding Rules** table, to add a new rule select the following options:
    a. For **Source**, select **User**.
    b. For **Destination**, select **Any**.
    c. For **Service**, select **UDP**. Enter the port range.
    d. For **Action**, select **Deny**.
    e. Click **Submit**.
    f. For the **Time range**, select **+** and enter the following for adding a new time range:
    - For **Name**, enter **working-hours**.
    - For **Type**, select **Periodic** and click +.
    - For **Start Day**, click **Weekday**.
    - For **Start Time**, enter **07:30.**
    - For **End Time**, enter **17:00**.
    - Click **Submit**.
11. Add another new rule for the guest-logon-access:
    a. For **Source**, select **Any**.
    b. For **Destination**, select **Any**.
    c. For **Service/app**, select **service**.

d. Select **svc-dhcp** for **Service alias**.

e. For **Action**, select **Permit.**

f. For **Time Range**, select **working-hours.**

g. Click **Submit**.

12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Creating Aliases

The following step defines an alias representing the public DNS server addresses. Once defined, you can use the alias for other rules and policies.

The following procedure describes how to create a destination alias:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Roles & Policies** > **Aliases** tab.
2. In **Network Aliases**, click **+**.
3. Select an **IP Version** from the drop-down list.
4. For **Name**, enter **Public DNS.**
5. For **Description**, enter a description of the destination within 128 characters.
6. Select **Invert** to specify that the inverse of the network addresses configured are used.
7. For **Items**, click **+**.
8. In the **Add New Destination Add New User Rule** window, for Rule Type, select **Host**. For IP Address, enter 64.151.103.120. **Click OK.**
9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Creating guest-logon-access policy

The following procedure describes how to create a guest-logon-access policy:

1. Login to the Mobility Master.
2. In the **Managed device** node hierarchy, navigate to the **Configuration > Roles & policies> Policies** page.
3. Select **+** to add the guest-logon-access policy.
4. For Policy Name, enter **guest-logon-access.**
5. For Policy Type, select **IPv4 Session**.
6. Click **Submit**.
7. Select the newly created **guest-logon-access** policy.
8. Click **+** under the **Policies > guest-logon-access** table.
9. In the **New Rule for guest-logon-access** popup, select **Access Control** option and click OK.
10. Under the **Roles > guest-logon-access > New forwarding Rules** table, to add a new rule select the following options:

    a. For **Source**, select **User**.

    b. For **Destination**, select **Alias**.

c. For **Destination Alias**, select **Public DNS.**

d. For **Service**, select **svc-dns.**

e. For **Action**, select **Source NAT**.

f. Under Time Range, select **working-hours**.

g. Click **Submit**.

11. Click **Pending Changes**.
12. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Creating an Auth-Guest-Access Policy

The following procedure describes how to configure the auth-guest-access policy:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Policies** tab.
3. Select **+** to create the policy.
4. For **Policy Name**, enter **auth-guest-access**.
5. For **Policy Type,** select **Session**.
6. Click **Submit**.
7. Select the newly created **auth-guest-access** policy.
8. Click **+** under the **Policies > auth-guest-access Roles** table.
9. In the **New Rule for auth-guest-access** popup, select **Access Control** option and click OK.
10. Under **auth-guest-access > New forwarding Rules**, to add a new rule select the following options:

    a. For **Source**, select **User**.

    b. For **Destination**, select **Any**.

    c. For **Service**, select **UDP**. Enter the port range.

    d. For **Action**, select **Deny**.

    e. Click **Submit**.

11. Repeat the steps 8 and 9 and in **auth-guest-access > New forwarding Rules**, select the following options to add another rule.

    a. For **Source**, select **Any**.

    b. For **Destination**, select **Any**.

    c. For **Service/app**, select **Service**.

    d. Select **svc-dhcp** for Service alias.

    e. For **Action**, select **Permit**.

    f. For **Time Range**, select **working-hours**.

    g. Click **Submit**.

12. Repeat the steps 8 and 9 and under **auth-guest-access > New forwarding Rules**, select the following options to add another rule.

a. For **Source**, select **User**.
b. For **Destination**, select **Alias**.
c. For **Destination Alias**, select **Public DNS** from the drop-down list.
d. For **Service/app**, select **Service**.
e. Select **svc-dns** for Service Alias.
f. For **Action**, select **Source NAT**.
g. For **Time Range**, select **working-hours**.
h. Click **Submit**.

13. Repeat steps 8 and 9 and under **auth-guest-access > New forwarding Rules**, select the following options to add another rule.
    a. For **Source**, select **User**.
    b. For **Destination**, select **Any**.
    c. For **Service/app**, select **Service**.
    d. Select **svc-http** for **Service alias**.
    e. For **Action**, select **Source NAT**.
    f. For **Time Range**, select **working-hours**.
    g. Click **Submit**.
14. Repeat the steps 8 and 9 and under **auth-guest-access > New forwarding Rules** table, select the following options to add another rule.
    a. For **Source**, select **User**.
    b. For **Destination**, select **Any**.
    c. For **Service/app**, select **service**.
    d. Select **svc-https** for **Service alias**.
    e. For **Action**, select **Source NAT**.
    f. For **Time Range**, select **working-hours**.
    g. Click **Submit**.
15. Click **Submit**.
16. Click **Pending Changes**.
17. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Creating an Block-Internal-Access Policy



It is recommended to first create a destination alias **Internalnetwork** and then create a block-internal-access policy. If the destination alias **Internalnetwork** is already created, the user can directly create a block - internal - access policy and skip the procedure given below.

The following procedure describes how to create a block-internal-access policy:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Roles & Policies** > **Aliases** tab.
2. In **Network Aliases**, click **+**.
3. Select an **IP Version** from the drop-down list.
4. For **Name**, enter **Internalnetwork**
5. For **Description**, enter a description of the destination within 128 characters.
6. Select **Invert** to specify that the inverse of the network addresses configured are used.
7. For **Items**, click **+**.
8. In the **Add New Destination Add New User Rule** window, for Rule Type, select **Network**. For IP Address, enter 10.0.0.0. For Network Mask or Range, enter 255.0.0.0. **Click OK.**
9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following procedure describes how to create the block-internal-access policy:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Policies** tab.
3. Select **+** to add a new policy.
4. For Policy Name, enter **block-internal-access**.
5. For Policy Type, select **Session**.
6. Click **Submit**.
7. Select the newly created **block-internal-access** policy.
8. Click **+** under the **Policies > block-internal-access** table.
9. In the **New Rule for block-internal-access** popup, select **Access Control** option and click OK.
10. Under the **block-internal-access > New forwarding Rules** table, to add a new rule select the following options:
    a. For Source, select **User**.
    b. For Destination, select **Alias**.
    c. For **Destination Alias**, select **Internalnetwork.**
    d. For **Service**, select **Any**.
    e. For **Action**, select **Deny**.
    f. Click **Submit**.
11. Click **Submit**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Creating a Drop-and-Log Policy

The following procedure describes how to create the drop-and-log policy:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Policies** tab.

3. Click **+** to add a new policy.
4. For **Policy Name**, enter **drop-and-log.**
5. For **Policy Type**, select **Session**.
6. Click **Submit**.
7. Select the newly created **drop-and-log** policy.
8. Click **+** under the **Policies > drop-and-log** table.
9. In the **New Rule for drop-and-log** popup, select **Access Control** option and click OK.
10. Under **drop-and-log > New forwarding Rules**, to add a new rule select the following options:
    a. For **Source**, select **User.**
    b. For **Destination**, select **Any**.
    c. For **Service**, select **Any.**
    d. For **Action**, select **Deny**.
    e. Select the **Log** checkbox from **Options**.
    f. Click **Submit**.
11. Click **Submit**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Creating a Guest Role

The following procedure describes how to create a guest role:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Roles** tab.
3. Click **+** to add a new role.
4. Enter **guest-logon** as a New Role.
5. Select the role name you just created and click **Show Advanced View**.
6. Click + under the **guest-logon role > Policies** tab.
7. In the **New Policy** popup, select the **Add an existing policy** option.
8. Select the policy name as **guest-logon** from the drop-down list
9. Click **Submit**.
10. Similarly, add block-internal-access policy for the role guest-logon.
11. Click **Submit**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Creating an Auth-Guest Role

The following procedure describes how to create the guest-logon role:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Roles** tab.
3. Click **+** to add a new role.

4. Enter **auth-guest** as a New Role.
5. Select the role name you just created and click **Show Advanced View**.
6. Click + under the **auth-guest role> Policies** tab.
7. In the **New Policy** pop-up, select the **Add an existing policy** option.
8. Select the policy name **cplogout** from the drop-down list
9. Click **Submit**.
10. Similarly, add guest-logon-access, block-internal-access, auth-guest-access, drop-and-log policies for the role **auth-guest**.
11. Click **Submit**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following set of CLI commands configures sample roles and policies:

Defining a Time Range

The following CLI command creates a time range:

```
(host) [md] (config) #time-range working-hours periodic
weekday 07:30 to 17:00
```

Creating Aliases

The following CLI commands create aliases:

```
(host) [md] (config) #netdestination "Internal Network"
network 10.0.0.0 255.0.0.0
network 172.16.0.0 255.255.0.0
network 192.168.0.0 255.255.0.0
(host)(config) #netdestination "Public DNS"
host 64.151.103.120
host 216.87.84.209
```

Creating a Guest-Logon-Access Policy

The following CLI commands create a guest-logon-access policy:

```
(host)(config) #ip access-list session guest-logon-access
user any udp 68 deny
any any svc-dhcp permit time-range working-hours
user alias "Public DNS" svc-dns src-nat time-range working-hours
```

Creating an Auth-Guest-Access Policy

The following CLI commands create an auth-guest-access policy:

```
(host) [md] (config) #ip access-list session auth-guest-access
user any udp 68 deny
any any svc-dhcp permit time-range working-hours
user alias "Public DNS" svc-dns src-nat time-range working-hours
user any svc-http src-nat time-range working-hours
user any svc-https src-nat time-range working-hours
```

Creating a Block-Internal-Access Policy

The following CLI command creates a block-internal-access policy:

```
(host) [md] (config) #ip access-list session block-internal-access
```

```
user alias "Internal Network" any deny
```

Creating a Drop-and-Log Policy

The following CLI command creates a drop-and-log policy:

```
(host) [md] (config) #ip access-list session drop-and-log
user any any deny log
```

Creating an Auth-Guest Role

The following CLI commands create an auth-guest role:

```
(host) [md] (config) #user-role auth-guest
(host) [md] (config-submode)#access-list session captiveportal
(host) [md] (config-submode)#access_list cplogout position 1
(host) [md] (config-submode)#access_list guest-logon-access position 2
(host) [md] (config-submode)#access_list block-internal-access position 3
(host) [md] (config-submode)#access_list auth-guest-access position 4
(host) [md] (config-submode)#access_list drop-and-log position 5
```

Configuring Guest VLANs

Guests using the WLAN are assigned to VLAN 900 and are given IP addresses via DHCP from the managed devices.

The following procedure describes how to configure a guest VLAN:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Interfaces > VLANs** tab.
3. Click + to add a new VLAN.
4. Enter the **VLAN name** as guest_vlan.
5. Enter **VLAN ID/Range** as 900.
6. Click **Submit**.
7. Select the VLAN name from the **VLANs** table and the **VLANs > <VLAN name>** table is displayed.
8. Click on the VLAN ID, 900 and enter the following:.
   a. For IPv4 Address, enter 192.168.200.20.
   b. Click **Submit**.
9. Navigate to the **Configuration > Services > DHCP Server** tab.
   a. Select **Enabled** for **IPV4 DHCP Server**.
   b. Click **+** under Pool Configuration.
   c. In the **Pool Name** field, enter **guestpool**.
   d. In the **Default Router** field, enter 192.168.200.20.
   e. In the **DNS Server** field, enter 64.151.103.120.
   f. In the **Lease hrs** field, enter 4 hours.
   g. Set the **Network IP address type** to Static.

h. In the **Network IP address** field, enter 192.168.200.0. In the **Network IP mask** field, enter 255.255.255.0.

i. Click **Submit.**

10. Click **Submit**.
11. Click **Pending Changes**.
12. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure a guest VLAN:

```
host) [mynode] #cd /md /<MAC_address>
(host) [<MAC_address>] (config)
(host) [<MAC_address>] (config) #vlan 900
(host) [<MAC_address>] (config) #interface vlan 900
(host) [<MAC_address>] (config) #ip address 192.168.200.20 255.255.255.0
(host) [<MAC_address>] (config) #ip dhcp pool "guestpool"
(host) [<MAC_address>] (config) #default-router 192.168.200.20
(host) [<MAC_address>] (config) #dns-server 64.151.103.120
(host) [<MAC_address>] (config) #lease 0 4 0
(host) [<MAC_address>] (config) #network 192.168.200.0 255.255.255.0
```

Configuring Captive Portal Authentication Profiles

In this section, you create an instance of the captive portal authentication profile and the AAA profile. For the captive portal authentication profile, you specify the previously-created **auth-guest** user role as the default user role for authenticated captive portal clients and the authentication server group (“Internal”).

The following procedure describes how to configure captive portal authentication:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > L3 Authentication** tab and select **Captive Portal Authentication**.
   a. Click + in the **Captive Portal Authentication Profile: New Profile** window to create a new Captive Portal Authentication Profile, enter **guestnet** as the **Profile Name** and click **Submit**.
   b. Select the captive portal authentication profile you just created.
   c. For Default Role, select **guest**.
   d. Select User Login.
   e. Uncheck **Guest Login**.
   f. Click **Submit**.
3. Select **Server Group** under the **guestnet** captive portal authentication profile you just created.
   a. Select **internal** from the **Server Group** drop-down list.
   b. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure captive portal authentication:

```
(host) [md] (config) #aaa authentication captive-portal guestnet
default-role auth-guest
```

```
user-logon
no guest-logon
server-group internal
```

The following section describes how to configure the user accounts, WLAN, AAA profile, and captive portal parameters:

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. You also need to include the predefined **captiveportal** policy, which directs clients to the captive portal, in the initial user role configuration.

The following procedure describes how to modify the guest-logon role:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Roles** tab.
3. Select the **guest-logon** role.
4. Select **Show Advanced View** in the **Roles > guest-logon** table.
5. Select the **More** tab.
6. Expand the **Authentication** accordion.
7. Select the captive portal authentication profile you just created from the **Captive Portal Profile** drop-down list, and click **Submit.**
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands modify the guest-logon role:

```
(host) [md] (config) #user-role guest-logon
(host) [md] (config-submode)#access-list session captiveportal
 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.

The following procedure describes how to configure the AAA profile:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > AAA Profiles** page.
3. Expand **AAA**. In the **AAA Profiles: New Profile**, click **+** to add a new profile. Enter **guestnet** for the name of the profile, then click **Submit**.
4. Select **guest-logon** from **Initial role** drop-down list.
5. Click **Submit.**
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures the AAA profile:

```
(host)[md](config) #aaa profile guestnet
initial-role guest-logon
```

Configuring the WLAN

In this section, you create the **guestnet** virtual AP profile for the WLAN. The **guestnet** virtual AP profile contains the SSID profile **guestnet** (which configures opensystem for the SSID) and the AAA profile **guestnet**.

The following procedure describes how to configure the guest WLAN:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
3. Under **All Profiles**, select **Wireless LAN**, then select **Virtual AP**.
4. To create a new virtual AP profile, Click **+** from the **Virtual AP profile: New Profile** pane. Enter the name for the virtual AP profile (for example, **guestnet**), and click **Submit**.
    a. In the Profile Details entry for the new virtual AP profile (**guestnet**), select **AAA profile** and then select the AAA profile you previously configured from the **AAA Profile** drop-down list and click **Submit**.
    b. In the Profile Details entry for the new virtual AP profile (**guestnet**), select **SSID** and then select SSID from the SSID profile drop-down list.
    c. Enter the name for the **ESSID** profile (for example, **guestnet**).
    d. For **Encryption**, select **opensystem.**
    e. At the bottom of the Profile Details page, click **Submit.**
5. Navigate to the **Configuration > AP Groups** page.
6. Select an AP Group and Click **WLANs** tab in the AP group window.
7. Click + under the WLANs tab and select the newly create virtual AP profile (guestnet) from the **Virtual-ap** drop-down list and click **Submit**.
8. Navigate to the **System > Profiles** tab. Select **Wireless LAN** and then select **Virtual AP**. Click on the new virtual AP name in the All Profiles list.
    a. Click **General** accordion and make sure **Virtual AP enable** is selected.
    b. For VLAN, enter the ID of the VLAN in which captive portal users are placed (for example, VLAN **900)**.
    c. Click **Submit.**
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the guest WLAN:

```
(host) [md] (config) #wlan ssid-profile guestnet
essid guestnet
opmode opensystem

(host) [md] (config) #aaa profile guestnet
initial-role guest-logon

(host) [md] (config) #wlan virtual-ap guestnet
vlan 900
```

```
aaa-profile guestnet
ssid-profile guestnet
```

Managing User Accounts

Temporary user accounts are created in the internal database on the Mobility Master. 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.

Configuring Captive Portal Configuration Parameters

Table 65 describes configuration parameters in the WebUI Captive Portal Authentication profile page.



In the CLI, you configure these options with the **aaa authentication captive-portal** commands.

Table 65: *Captive Portal Authentication Profile Parameters*

Parameter	Description
Default Role	Role assigned to the Captive Portal user upon login. When both user and guest logon are enabled, the default role applies to the user logon; users logging in using the guest interface are assigned the guest role. Default: guest
Default Guest Role	Role assigned to guest. Default: guest
Redirect Pause	Time, in seconds, that the system remains in the initial welcome page before redirecting the user to the final web URL. If set to 0, the welcome page displays until the user clicks on the indicated link. Default: 10 seconds
User Login	Enables Captive Portal with authentication of user credentials. Default: Enabled
Guest Login	Enables Captive Portal logon without authentication. Default: Disabled
Logout popout window	Enables a pop-up window with the Logout link for the user to logout after logon. If this is disabled, the user remains logged in until the user timeout period has elapsed or the station reloads. Default: Enabled
Use HTTP for authentication	Use HTTP protocol on redirection to the Captive Portal page. If you use this option, modify the captive portal policy to allow HTTP traffic. Default: disabled (HTTPS is used)
Logon wait minimum wait	Minimum time, in seconds, the user will have to wait for the logon page to pop up if the CPU load is high. This works in conjunction with the Logon wait CPU utilization threshold parameter. Default: 5 seconds
Logon wait maximum wait	Configure parameters for the logon wait interval Default: 10 seconds
Logon wait CPU utilization threshold	CPU utilization percentage above which the Logon wait interval is applied when presenting the user with the logon page.

Parameter	Description
	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 FQDN on the login page. The FQDNs shown are specified when configuring individual servers for the server group used with captive portal authentication. Default: Disabled
Authentication Protocol	Select the PAP, CHAP or MS-CHAPv2 authentication protocol. **NOTE:** Do not use the CHAP = option unless instructed to do so by anAruba representative.
Login Page	URL of the page that appears before logon. This can be set to any URL. Default: `/cgi-bin/login?cmd=authenticate` or `/cgi-bin/login?cmd=login`
Welcome Page	URL of the page that appears after logon and before redirection to the web URL. This can be set to any URL. Default: /auth/welcome.html
Show Welcome Page	Displays the configured welcome page before the user is redirected to their original URL. If this option is disabled, users are redirected to the web URL immediately after they log in. Default: Enabled
Proxy Server Configuration	To configure proxy details for captive portal authentication. **NOTE:** User cannot configure this setting.
Add switch IP address in redirection URL	Sends the IP address of the managed device in the redirection URL when external captive portal servers are used. An external captive portal server can determine the managed devices from which a request originated by parsing the 'switchip' variable in the URL. Default: Disabled
Adding User VLAN in the redirection URL	Sends the user VLAN ID in the redirection URL when external captive portal servers are used.
Adding AP's MAC address in redirection URL	AP's MAC address is added in the redirection URL when external captive portal servers are used. Default: Disabled
Add a controller interface in the redirection URL	Sends the interface IP address of the managed device in the redirection URL when external captive portal servers are used. An external captive portal server can determine the managed devices from which a request originated by parsing the 'switchip' variable in the URL.
Allow only one active user session	Allows only one active user session at a time. Default: Disabled

Parameter	Description
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 a PEFNG license.
Black List	To add a netdestination to the captive portal blacklist, enter the destination host or subnet, then click **Add**. The netdestination will be added to the blacklist. To remove a netdestination from the blacklist, select it in the blacklist field, then click **Delete**. If you have not yet defined a netdestination, use the CLI command **netdestination** to define a destination host or subnet before you add it to the blacklist.
Show Acceptable Use Policy Page	Show the acceptable use policy page before the logon page. Default: Disabled
User idle timeout	The user idle timeout value for this profile. Specify the idle timeout value for the client in seconds. Valid range is 30-15300 in multiples of 30 seconds. Enabling this option overrides the global settings configured in the AAA timers. If this is disabled, the global settings are used.
Redirect URL	URL to which an authenticated user will be directed. This parameter must be an absolute URL that begins with either **http://** or **https://**.
Bypass Apple and Android Captive Network Assistant	Enabling this knob will bypass Apple CNA on iOS devices like iPad, iPhone, iPod and on Android devices. The user needs to perform Captive Portal authentication from the browser.
URL Hash Key	If a redirection URL is defined, enter a URL Hash Key to hash the redirect URL using the specified key. This parameter enhances security for the ClearPass Guest login URL so that ClearPass Policy Manager can trust and ensure that the client MAC address in the redirect URL has not been tampered with by anyone. Default: Disabled.

Related Topics

Enabling Guest Provisioning

Configuring Captive Portal in the Base Operating System

Configuring Captive Portal with a PEFNG License

Enabling Optional Captive Portal Configuration

You can configure optional captive portal pages by using the WebUI or the CLI.

This section describes the following topics:

- Uploading Captive Portal Pages by SSID Association on page 349
- Changing the Protocol to HTTP on page 350
- Configuring Redirection to a Proxy Server on page 351
- Redirecting Clients on Different VLANs on page 352
- Web Client Configuration with Proxy Script on page 353

Uploading Captive Portal Pages by SSID Association

You can upload custom login pages for captive portal into the managed device through the WebUI. 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 managed device, 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 66.

Table 66: *Captive Portal login Pages*

Entity	Engineering	Business	Faculty
Captive portal login page	eng-login.html	bus-login.html	fac-login.html
Captive portal user role	eng-user	bus-user	fac-user
Captive portal authentication profile	eng-cp (Specify eng-login.html and eng-user)	bus-cp (Specify bus-login.html and bus-user)	fac-cp (Specify 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.

The following procedure describes how to change the protocol to HTTP:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, edit the captive portal authentication profile by navigating to the **Configuration > Authentication > L3 Authentication** tab.

3. Select a captive portal profile, enable the **Use HTTP for authentication** check box and click **Submit**.
4. (For captive portal with role-based access only) Edit the **captive portal** policy by navigating to the **Configuration > Roles & Policies > Policies** tab.
   a. Select the policy for which you want to add or delete a new rule.
   b. Click **+** in the **Policy > <name of the policy> Rules** table. Select a **Rule type** and click **Ok**.
   c. Add a new rule with the following values:
   - **Source** is User.
   - **Destination** is the mswitch alias.
   - **Service** is svc-http.
   - **Action** is dst-nat.

   d. Click **Submit.**
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands change the protocol to HTTP:

```
(host) [md] (config) #aaa authentication captive-portal profile
protocol-http

(For captive portal with role-based access only)
(host) [md] (config) #ip access-list session captiveportal
no user alias mswitch svc-https dst-nat
user alias mswitch svc-http dst-nat
user any svc-http dst-nat 8080
user any svc-https dst-nat 8081
```

Configuring Redirection to a Proxy Server

You can configure captive portal to work with proxy Web servers. When proxy Web servers are used, browser proxy server settings for end users are configured for the IP address and TCP port of the proxy server. When the user opens a Web browser, the HTTP or HTTPS connection request must be redirected from the proxy server to the captive portal on the managed devices.

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 IP address and TCP port of the proxy server.
- (For captive portal with role-based access) Modify the **captiveportal** policy to have traffic for the port destination of the proxy server with NAT applied to port 8088 on the managed device.

The base operating system automatically modifies the implicit ACL *captive-portal-profile*.

The following sections describe how use the WebUI and CLI to configure the captive portal with a proxy server.



When HTTPS traffic is redirected from a proxy server to the managed device, the users browser will display a warning that the subject name on the certificate does not match the hostname to which the user is connecting.

The following procedure describes how to redirect proxy server traffic:

1. Login to the Mobility Master.
2. For captive portal with Aruba base operating system, in the **Managed Network** node hierarchy, edit the captive portal authentication profile by navigating to the **Configuration > Authentication > L3 Authentication** page.
    a. Select a captive portal profile and enter the IP address and port for the proxy server.
    b. Click **Submit**.
3. For captive portal with role-based access, edit the **captiveportal** policy by navigating to the **Configuration > Roles and Policies > Policies** tab.
4. Select the policy you want to edit.
5. Click **+** in the **Policy > <name of the policy> Rules** table. Select a **Rule type** and click **Ok**.
6. 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.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands redirect proxy server traffic:

For captive portal with Aruba base operating system:

```
(host) [md] (config) #aaa authentication captive-portal profile
proxy host ipaddr port port
For captive portal with role-based access:
(host) [md] (config) #ip access-list session captiveportal
user alias mswitch svc-https permit
user any tcp port dst-nat 8088
user any svc-http dst-nat 8080
user any svc-https dst-nat 8081
```

Redirecting Clients on Different VLANs

You can redirect wireless clients that are on different VLANs (from the IP address of the managed device) to the captive portal on the managed device. 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 managed device interface.
    b. Modify the rule set to allow HTTPS to the new alias instead of the mswitch alias.



In the base operating system, the implicit ACL *captive-portal-profile* is automatically modified.

This example shows how to use the command-line interface to create a network destination called cp-redirect and use that in the captive portal policy:

```
(host) [md] (config     ) #ip cp-redirect-address ipaddr
```

For captive portal with PEFNG license:

```
(host) [md] (config) #netdestination cp-redirect
(host) [md] (config-submode)#ip access-list session captiveportal
user alias cp-redirect svc-https permit
user any svc-http dst-nat 8080
user any svc-https dst-nat 8081
```

Web Client Configuration with Proxy Script

If the web client proxy configuration is distributed through a proxy script (a `.pac` file), you need to configure the **captiveportal** policy to allow the client to download the file. Note that in order modify the captiveportal policy, you must have the PEFNG license installed in the managed device.

The following procedure describes how to allow clients to download proxy script:

1. Login to the Mobility Master.
2. Edit the **captiveportal** policy by navigating to the **Configuration > Roles & Policies > Policies** tab in the **Managed Network** node hierarchy.
3. Select the policy you want to edit.
4. Click **+** in the **Policy > <name of the policy> Rules** table. Select a **Rule type** and click **Ok**.
5. 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**.
6. Click **Submit** to add the rule.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands allow clients to download proxy script:

```
(host) [md] (config) #ip access-list session captiveportal
user alias mswitch svc-https permit
user any tcp port dst-nat 8088
user host ipaddr svc-https permit
user any svc-http dst-nat 8080
user any svc-https dst-nat 8081
```

Related Topics

Personalizing the Captive Portal Page

Creating and Installing an Internal Captive Portal

Personalizing the Captive Portal Page

The following can be personalized on the captive portal page:

- Captive portal background
- Welcome text
- Acceptance Use Policy

Starting with ArubaOS 8.0.0.0, Reply-Message that is returned by RADIUS server for a Captive Portal Authentication can be customized using the Standard RADIUS attribute **reply-Message** VSA.

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.



Captive Portal profile have few configurations which are confined only to WebUI and there are no command line interface commands to perform some of the actions like uploading custom login or Welcome page, background images, logos, Acceptable Usage Policy texts, and so on.

This section describes the following topics:

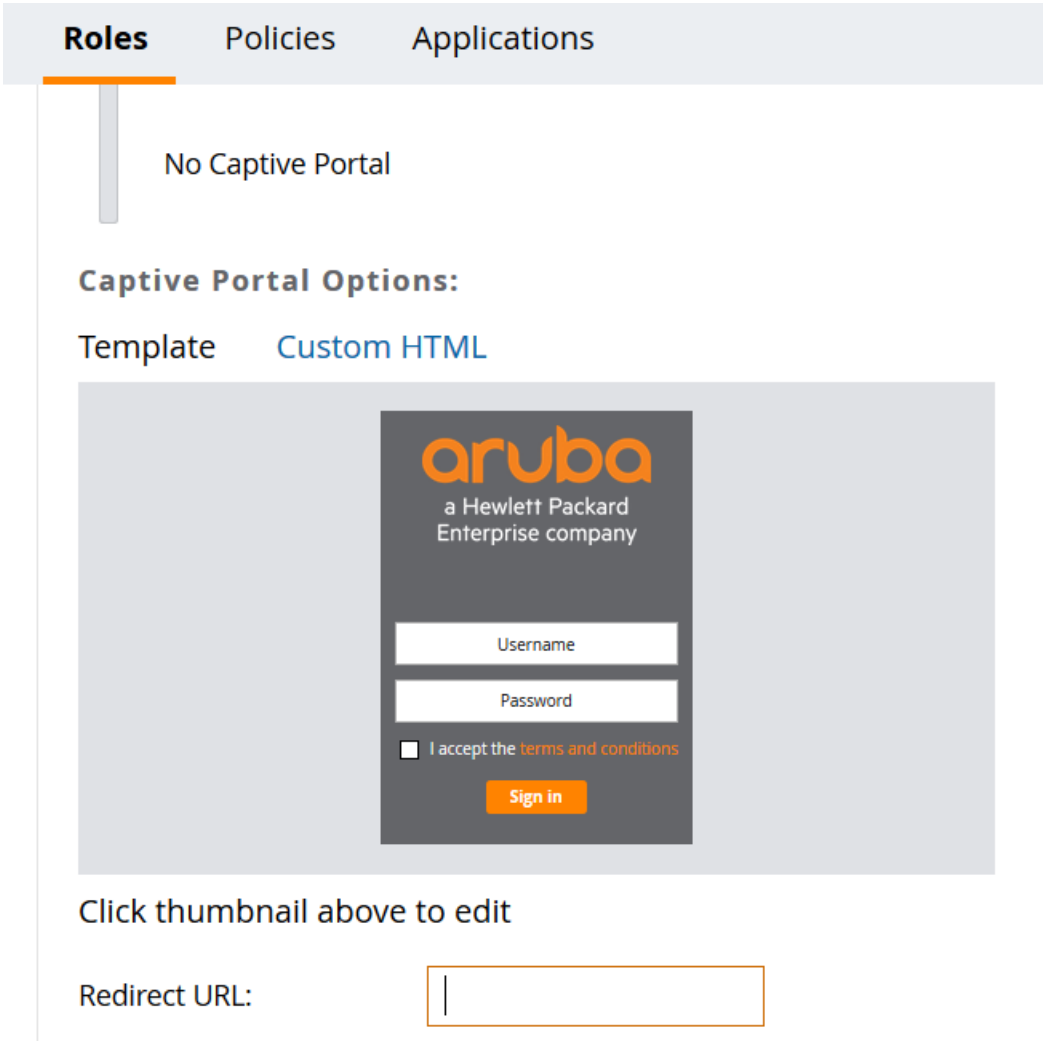
- Creating your Own Web Pages and Install them on page 354
- Customizing the Captive Portal Page for a Role on page 357

Creating your Own Web Pages and Install them

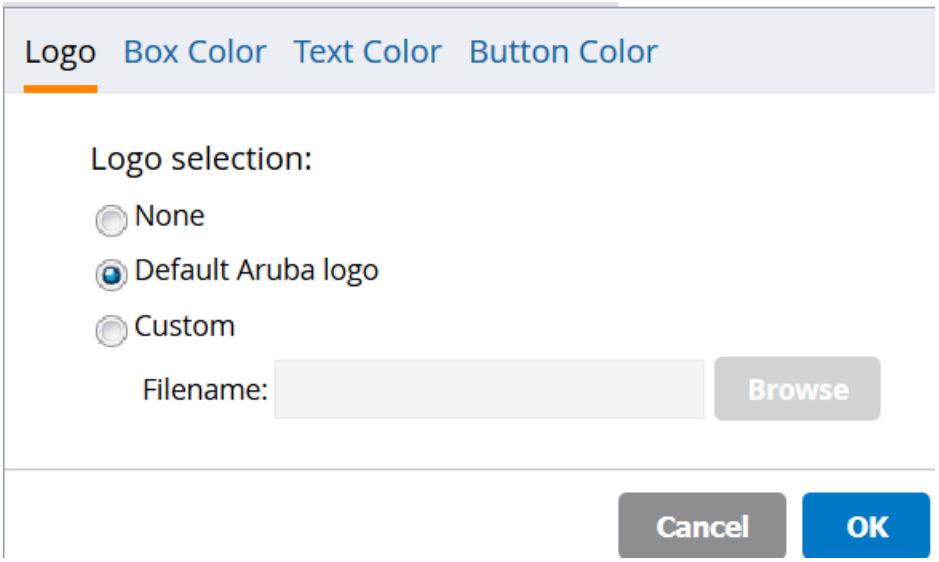
The following procedure describes how to create your own web pages and install them in the managed device:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies> Roles** tab.
3. Select a role and click **Show Advanced View** .

4. Click on the **Captive Portal** tab. Click **Internal captive portal with authentication** option.

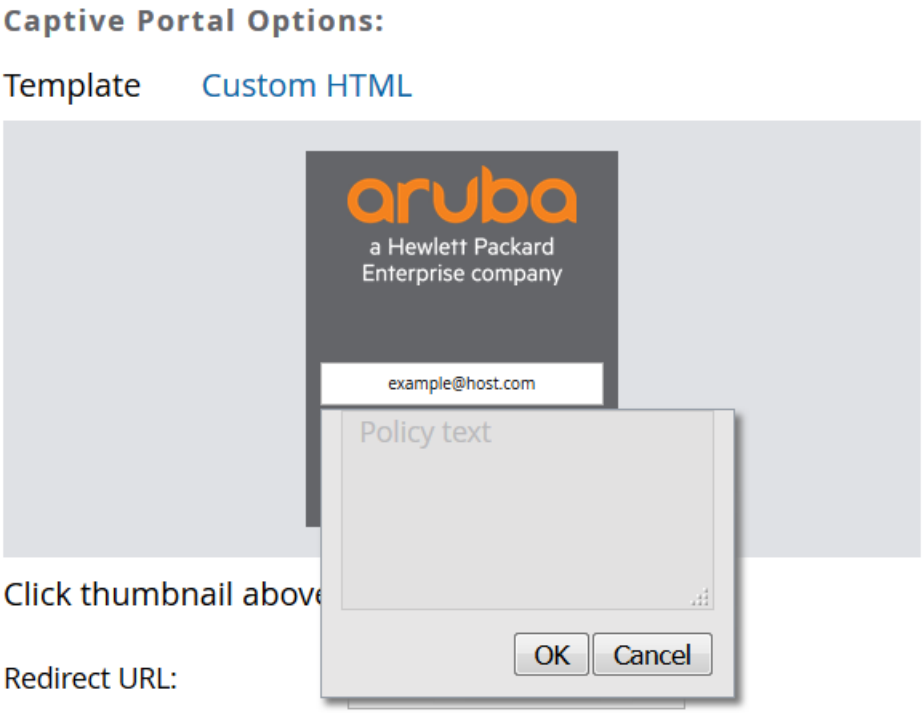


5. Click on the Thumbnail to edit the Templates. You can edit the logo, Box Color, Text Color, and Button Color. Click **Preview** the view the changes.



6. Click **Submit** to save the changes.

7. You can also change the AUP text using templates link in the **Captive Portal Options** window. Click on **I accept the Terms and Conditions** option and the window to add **Policy Text** is displayed.

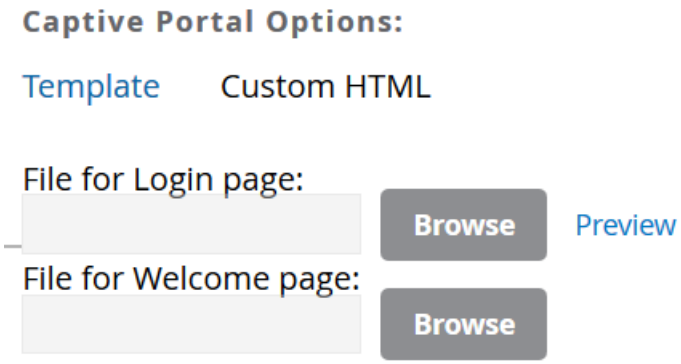


8. Click **Submit** to save the changes.

NOTE

When the Terms and conditions link is clicked, the AUP text is displayed only if the AUP text was previously entered.

9. To upload Login or Welcome page, perform the following steps using Custom HTML link in the **Captive Portal Options** window:
   a. Click on Custom HTML link.



   b. To change the login page, browse for the file through the **File for Login Page** option.
   c. To change the Welcome page, browse for the file through the **File for Welcome Page** option.

d. Before submitting the changes, ensure that the changes are accurate by clicking the **Preview** option.

10. . Similarly, you can customize the page for Internal Captive Portal with email registration and for Internal Captive portal, no auth or registration.
11. Click **Submit**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Customizing the Captive Portal Page for a Role

Captive Portal page can also be customized for a particular user role.

The following procedure describes how to customize the Captive Portal page for a role:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to **Configuration > Roles & Policies > Roles** tab.
3. Select the role you want to customize the Captive Portal page for.
4. Select **Show Advanced View**.
5. Click **Captive Portal** tab.
6. Click **Internal captive portal with email registration** option and the **Captive Portal Options** are displayed.
7. You can also customize the Captive Portal page while creating the virtual AP with option as Guest in the **Configuration > WLAN** options.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Related Topics

Enabling Optional Captive Portal Configuration

Creating and Installing an Internal Captive Portal

Creating and Installing an Internal Captive Portal

If you do not wish to customize the default captive portal page, you can create and install a new internal captive portal page.

This section describes the following topics:

- Creating a New Internal Web Page on page 357
- Installing a New Captive Portal Page on page 359
- Displaying Authentication Error Messages on page 360
- Configuring Localization on page 360
- Customizing the Welcome Page on page 364
- Customizing the Pop-Up box on page 365
- Customizing the Logged Out Box on page 366

Creating a New Internal Web Page

In addition to customizing the default captive portal page, you can also create your own internal web page. A custom web page must include an authentication form to authenticate a user. The authentication form can include any of the following variables listed in Table 67:

Table 67: *Web Page Authentication Variables*

Variable	Description
user	(Required)
password	(Required)
FQDN	The fully-qualified domain name (this is dependent on the setting of the managed device 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://<managed device_IP>/cgi-bin/login.

You can construct an authentication form using the following HTML:

```
<form name="form1" method="post" action="/cgi-bin/login">
...
</FORM>
A recommended option for the <FORM> element is:
autocomplete="off"
For example
<form name="form1" method="post" action="/cgi-bin/login">
<div id="header">
<h1 id="logo"><a href="#"> </a></h1>
</div>
<input type="hidden" id="email" name="email" type="text" value="guest@abc.com"
class="text" accesskey="e" />
<input type="hidden" name="cmd" value="authenticate" />
<input type="submit" name="Login" value="ACCEPT" class="button" />
</form>
```

This option prevents Internet Explorer from caching the form inputs. The form variables are input using any form control method available such as INPUT, SELECT, TEXTAREA, and BUTTON. Example HTML code follows.

Username Example

Minimal:

```
        <INPUT type="text" name="user">
```

Recommended Options:

```
    accesskey="u"   Sets the keyboard shortcut to 'u'
    SIZE="25          "Sets the size of the input box to 25
    VALUE=       ""Ensures no default value
```

Password Example

Minimal:

```
    <INPUT type="password" name="password">
```

Recommended Options:

```
    accesskey="p"   Sets the keyboard shortcut to 'p'
    SIZE="25          "Sets the size of the input box to 25
```

```
    VALUE=      ""Ensures no default value
```

FQDN Example

Minimal:

```
<SELECT name=fqdn>
    <OPTION value="fqdn1" SELECTED>
    <OPTION value="fqdn2">
</SELECT>
```

Recommended Options:

```
    None
```

Finally, an HTML also requires an input button:

```
<INPUT type="submit">
```

Basic HTML Example

```
<HTML>
 <HEAD>
 </HEAD>
 <BODY>
  <FORM method="post" autocomplete="off" ACTION="/cgi-bin/login">

  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

The following procedure describes how to install the captive portal page by using the Maintenance function:

1. Login to the Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Roles** tab and select a role and click **Show Advanced View** in **Roles > <rolename>** table.
3. Click **Captive Portal**. Click **Internal captive portal with authentication**.
4. Click on **Custom HTML**.
5. To change the login page, browse for the file through the **File for Login Page** option.
6. To change the Welcome page, browse for the file through the **File for Welcome Page** option.

This page lets you upload your own files to the managed device. There are different page types that you can choose:

- Captive Portal Login (top level): This type uploads the file into the managed device and sets the captive portal page to reference the file that you are uploading. Use with caution on a production managed device 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.

Uploaded files can be referenced using:

```
https://<managed device_IP>/upload/custom/<CP-Profile-Name>/<file>
```

Displaying Authentication Error Messages

This section contains a script that performs the following tasks:

- When the user is redirected to the main captive portal login when there is authentication failure, the redirect URL includes a query parameter "errmsg" which java script can extract and display.
- Store the originally requested URL in a cookie so that once the user has authenticated, they are automatically redirected to its original page.

```
<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>
```

Configuring Localization

The ability to customize the internal captive portal provides you with a very flexible interface to the Aruba 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 Aruba internal captive portal system.

1. Customize the configurable parts of the captive portal settings to your liking. To do this, navigate to the **Configuration > Roles & Policies page.** Edit **Role** and select **Show Advanced tab** and click **Captive Portal.**

   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 Guest login is enabled or disabled as necessary by selecting the right option from path **Configuration > Roles & Policies page. Edit Role and select Show Advanced tab and click Captive Portal** to create or edit the captive portal profile.

2. Click **Submit** and then click on **Preview.** 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 **Preview** 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 on page 361, 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");}
</script>
</head>
```

c. Fix references: If you have used the built-in preferences, you will need to update the reference for the logo image and the CSS style sheet.

   To update the CSS reference, search the text for "<link href" and update the reference to include "/auth/" in front of the reference. The original link should look similar to the following:

```
<link href="default1/styles.css" rel="stylesheet" media="screen" type="text/css" />
```

   This should be replaced with a link like the following:

```
<link href="/auth/default1/styles.css" rel="stylesheet" media="screen" type="text/css" />
```

   The easiest way to update the image reference is to search for "src" using your text editor and updating the reference to include "/auth/" in front of the image file. The original link should look similar to the following:

```
<img src="default1/logo.gif"/>
```

   This should be replaced with a link like this:

```
<img src="/auth/default1/logo.gif"/>
```

d. Insert javascript to handle error cases:

   When the managed device detects an error situation, it will pass the user's page a variable called "errmsg" with a value of what the error is in English. Currently, only "Authentication Failed" is supported as a valid error message.

   To localize the authentication failure message, replace the following text (it is just a few lines below the <body> tag):

```
<div id="errorbox" style="display: none;">
</div>
```

   with the script below. You will need to translate the "Authentication Failed" error message into your local language and add it into the script below where it states: localized_msg="...":

```
<script>
{
 var q = window.location.search;
 var errmsg = null;
 if (q && q.length > 1) {
  q = q.substring(1).split(/[=&]/);
```

```
    for (var i = 0; i < q.length - 1; i += 2) {
      if (q[i] == "errmsg") {
        errmsg = unescape(q[i + 1]);
        break;
      }
    }
  }

  if (errmsg && errmsg.length > 0) {
    switch(errmsg) {
    case "Authentication Failed":
       localized_msg="Authentication Failed";
       break;
    default:
     localised_msg=errmsg;
     break;
    }
    errmsg = "<div id='errorbox'>\n" + localised_msg + "\n</div>\n";
    document.write(errmsg);
  };
}
</script>
```

e. Translate the web page text. Once you have made the changes as above, you only need to translate the rest of the text that appears on the page. The exact text that appears will depend on the managed device 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 Mobility Master using the navigation provided in the Installing a New Captive Portal Page section.

   Choose the captive portal profile from the drop-down list. 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 managed device 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 managed device in order to view the page again.

6. Finally, it is possible to customize the welcome page on the managed device, 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 managed device.

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 managed device, the user needs to create their own web page (using the charset meta attribute in step 4 above). Upload this page to the designated managed device in the same manner as uploading the captive portal login page under "**Configuration > Management > Captive Portal > Upload Custom Login Pages.** For Page Type, select “Captive Portal Welcome Page”.

Any required client side script (CSS) and media files can also be uploaded using the “Content” Page Type, however file space is limited (use the CLI command **show storage** to see available space). Remember to leave ample room for system files.



The "Registered User" and "Guest User" sections of the login page are implemented as graphics files, referenced by the default CSS styles. In order to change these, you will need to create new graphic files, download the CSS file, edit the reference to the graphics files, change the style reference in your index file and then upload all files as "content" to the managed device.

Customizing the Welcome Page

Once a user is authenticated by the managed device, a Welcome page is launched.

You can customize this welcome page by building your own HTML page and uploading it to the Mobility Master. You upload it to the Mobility Master using the navigation provided in the Installing a New Captive Portal Page section. This file is stored in a directory called "/upload/" on the Mobility Master 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 is shown below. The part in red will redirect the user to the web page you originally set up. 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="/cp/logout">
            <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="/cp/logout"> 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 /cp/logout Once a user accesses this URL then the managed device 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 managed device using the navigation provided in the Installing a New Captive Portal Page section.

Once you have completed your HTML, then you must get the clients to create the pop-up box once they have logged into the managed device. This is done by inserting the following code into your welcome page text and re-uploading the welcome page text to your managed device.

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 /cp/logout file on the managed device and so it is hidden in the html page here in order to get the client to access this page and for the managed device 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 /cp/logout. Either method should allow the client to logout from the managed device.

Everything else can be customized.

```
<html>
<body bgcolor=white text=000000>

<iframe src='/cp/logout' width=0 height=0 frameborder=0><img src=/cp/logout 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 "/cp/logout" with your own file that you upload to the managed device. 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>
```

Related Topics

Enabling Captive Portal Enhancements

Creating Walled Garden Access

Creating Walled Garden Access

On the Internet, a walled garden typically controls a user access to web content and services. The walled garden directs the user navigation within particular areas to allow access to a selection of websites or prevent access to other websites.



The Walled Garden feature can be used with the PEFNG or PEFV licenses.

Walled garden access is needed when an external or internal captive portal is used. A common example could be a hotel environment where unauthenticated users are allowed to navigate to a designated login page (for example, a hotel website) and all its contents.

Users who do not sign up for Internet service can view “allowed” websites (typically hotel property websites). The website names must be DNS-based (not IP address based) and support the option to define wildcards.

HTTP or HTTPS proxy does not work when walled garden is implemented as a user-role using domain name ACL. For example, **user alias example.com any permit**.

When a user attempts to navigate to other websites not configured in the white list walled garden profile, the user is redirected back to the login page. In addition, the black listed walled garden profile is configured to explicitly block navigation to websites from unauthenticated users.

The following CLI example configures a destination named Mywhite-list and adds the domain names, example.com and example.net to that destination using the CLI. It then adds the destination name Mywhite-list (which contains the allowed domain names example.com and example.net) to the white list.

```
(host) [md] (config)# netdestination "Mywhite-list"
(host) [md] (config)#name example.com
(host) [md] (config)#name example.net
```

```
(host) [md] (config) #aaa authentication captive-portal default
 (host) [md] (Captive Portal Authentication Profile "default")#white-list Mywhite-list
```

The following procedure describes how to configure a walled garden access:

1. Login to the Mobility Master.
2. In the **Managed Network**node hierarchy, navigate to **Configuration** > **Roles and Policies** > **Policies** tab.
3. Click **+** to add a new policy.
4. Enter **Policy Name** and set the **Policy Type** to Session.
5. Select the newly created policy name and Click + in **Policy <Name of the policy> Rules** to add a new rule.
6. Select Access Control as the **Rule Type** and click **OK**.
7. In the **New forwarding Rule** window:
   a. Select the IP version of the managed device, IPv4 or IPv6, from the **IP Version** drop-down list.
   b. Select the destination as **Alias**.
   c. Select the Destination Alias as **Mywhite-list**.
8. Click **Submit.**
9. Navigate to **Configuration > Authentication > L3 Authentication.**
10. Select **Captive Portal Authentication Profile and select a profile.**
11. To allow users to access a domain, enter the destination name that contains the allowed domain names in the **White List** field. This stops unauthenticated users from viewing specific domains such as a hotel website.

    A rule in the white list must explicitly permit a traffic session before it is forwarded to the managed device. The last rule in the white list denies everything else.

12. To deny users access to a domain, enter the destination name that contains prohibited domain names in the **Black List** field. This prevents unauthenticated users from viewing specific websites.
13. Click **Submit**.
14. Click **Pending Changes**.
15. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Related Topics

Enabling Captive Portal Enhancements

Enabling Captive Portal Enhancements

ArubaOS introduces the following enhancements in Captive Portal:

- Location information such as AP name and AP group name have been included in the Captive Portal redirect URL. The following example shows a Captive Portal redirect URL that contains the AP name and the AP group name:

https://securelogin.example.com/cgi-bin/login?cmd=login&mac=00:24:d7:ed:84:14&ip=10.15.104.13&essid=example-test-tunnel&apname=ap135&apgroup=example&url=http%3A%2F%2Fwww%2Eespncricinfo%2Ecom%2F

- A new option **redirect-url** is introduced in the Captive Portal Authentication profile which allows you to redirect the users to a specific URL after the authentication is complete.
- Captive Portal Login URL length has been increased from 256 characters to 2048 characters.
- Support for "?" (question mark) inside the Captive Portal login URL has been added.
- A new field, **description** has been introduced in the **netdestination** and **netdestination6** commands to provide a description about the netdestination up to 128 characters long.
- Support for configuring Whitelist in Captive Portal has been introduced.

The Captive Portal enhancements are available on Tunnel and Split-Tunnel forwarding modes.

The following section describes the various enhancements in Captive Portal:

Configuring the Redirect-URL

The following CLI commands configure the Captive Portal redirect URL:

```
(host) [md] (config) # aaa authentication captive-portal REDIRECT
(host) [md] (Captive Portal Authentication Profile "REDIRECT") #redirect-url
<absolute-URL>
Example:
(host) [md] (config) # aaa authentication captive-portal REDIRECT
(host) [md] (Captive Portal Authentication Profile "REDIRECT") #redirect-url
https://test-login.php
```

Configuring the Login URL

The following CLI commands configure a Captive Portal login URL up to 2048 characters:

```
(host) [md] (config) # aaa authentication captive-portal LOGIN
(host) [md] (Captive Portal Authentication Profile "LOGIN")#login-page
"https://clearpass-dev1.dev.arubademo.net/guest/aos8_self-reg.php?_browser=1"
```



You can configure the login URL with "?" (question mark) character in it provided the URL containing the question mark is within the double quotes.

Defining Netdestination Descriptions

You can provide a description (up to 128 characters) for the netdestination using the CLI.

The following CLI commands provide description for an IPv4 netdestination:

```
(host) [md] (config) #netdestination Local-Server
(host) [md] (config-dest) #description "This is a local server for IPv4 client
registration"
The following CLI commands provide description for an IPv6 netdestination:
(host) [md] (config) #netdestination6 Local-Server6
(host) [md] (config-dest) #description "This is a local server for IPv6 client
registration"
```

The following CLI command displays the details of the specified IPv4 netdestination in the managed device:

```
(host) (config-dest)#show netdestination Local-Server

Name: Local-Server
Description: This is a local server for IPv4 client registration
Position  Type  IP addr  Mask-Len/Range
--------  ----  -------  --------------
1         name  0.0.0.1   yahoomail
2         name  0.0.0.2   mycorp
3         name  0.0.0.3   cricinfo
```

The following CLI command displays the details of the specified IPv6 netdestination in the managed device:

```
(host) (config-dest) #show netdestination Local-Server6

Name: Local-Server6
Description: This is a local server for IPv6 client registration
------------------------------------------------------------------------------
Position  Type  IP addr   Mask-Len/Range
--------  ----  -------   --------------
1         name  ::9        yahoomail
2         name  ::a        mycorp
3         name  ::b        cricinfo
```

Configuring a Whitelist

You can now configure a whitelist in Captive Portal using the CLI.

This section describes the following topics:

Configuring the Netdestination for a Whitelist:

The following CLI commands configure a netdestination alias for Whitelist:

```
(host) [md] (config) #netdestination whitelist
(host) [md] (config-dest) #description guest_whitelist
(host) [md] (config-dest) #name mycorp
```

Associating a Whitelist to Captive Portal Profile

The following CLI commands associate a whitelist to the Captive profile:

```
(host) [md] (config) #aaa authentication captive-portal CP_Profile
(host) [md] (Captive Portal Authentication Profile "CP_Profile”) #white-list whitelist
```

Applying a Captive Portal Profile to a User-Role

The following CLI commands apply the Captive Portal profile to a user-role:

```
(host) [md] (config) # user-role guest_role
(host) [md] (config-submode) #access_list logon-control
(host) [md] (config-submode) #access_list captiveportal
(host) [md] (config-submode) #captive-portal CP_Profile
```

Verifying a Whitelist Configuration

The following CLI command verifies the whitelist alias in the managed device:

```
(host) (config) #show netdestination whitelist

whitelist Description: guest_whitelist
--------------------------------------
Position  Type  IP addr   Mask-Len/Range
--------  ----  -------   --------------
1         name  0.0.0.6   mycorp
```

Verifying a Captive Portal Profile Linked to a Whitelist

The following CLI command verifies the Captive Portal profile linked to the whitelist in the managed device:

```
(host) (config) #show aaa authentication captive-portal CP_Profile

Captive Portal Authentication Profile "CP_Profile"
--------------------------------------------------------------------
Parameter                                          Value
---------                                          -----
Default Role                                       guest
Default Guest Role                                 guest
Server Group                                       default
Redirect Pause                                     10 sec
User Login                                         Enabled
Guest Login                                        Disabled
Logout popup window                                Enabled
Use HTTP for authentication                        Disabled
Logon wait minimum wait                            5 sec
Logon wait maximum wait                            10 sec
logon wait CPU utilization threshold               60 %
Max Authentication failures                        0
Show FQDN                                          Disabled
Use CHAP (non-standard)                            Disabled
Login page                                         /auth/index.html
Welcome page                                       /auth/welcome.html
Show Welcome Page                                  Yes
Add switch IP address in the redirection URL       Disabled
Adding user vlan in redirection URL                Disabled
Add a controller interface in the redirection URL  N/A
Allow only one active user session                 Disabled
White List                                         whitelist
Black List                                         N/A
Show the acceptable use policy page                Disabled
Redirect URL                                       N/A
```

Verifying Dynamic ACLs for a Whitelist

The following CLI command verifies the dynamically created ACLs for the whitelist in the managed device:

```
(host) (config)#show rights guest_role

Derived Role = 'guest_role'
Up BW:No Limit   Down BW:No Limit
L2TP Pool = default-l2tp-pool
PPTP Pool = default-pptp-pool
Periodic reauthentication: Disabled
ACL Number = 79/0
Max Sessions = 65535
Captive Portal profile = CP_Profile

access-list List
----------------
Position  Name                                        Location
--------  ----                                        --------
1         CP_Profile_list_operations
2         logon-control
3         captiveportal
CP_Profile_list_operations
--------------------------------------
Priority  Source  Destination  Service    Action  TimeRange  Log  Expired  Queue  TOS
8021P  Blacklist  Mirror  DisScan  ClassifyMedia  IPv4/6
--------  ------  -----------  -------    ------  ---------  ---  -------  -----  ---
-----  ---------  ------  -------  -------------  ------
1         user    whitelist    svc-http   permit                          Low
                                                4
2         user    whitelist    svc-https  permit                          Low
                                                4
logon-control
-------------
Priority  Source  Destination  Service   Action  TimeRange  Log  Expired  Queue  TOS
8021P  Blacklist  Mirror  DisScan  ClassifyMedia  IPv4/6
--------  ------  -----------  -------   ------  ---------  ---  -------  -----  ---
-----  ---------  ------  -------  -------------  ------
1         user    any          udp 68    deny                            Low
                                                4
2         any     any          svc-icmp  permit                          Low
                                                4
3         any     any          svc-dns   permit                          Low
                                                4
4         any     any          svc-dhcp  permit                          Low
                                                4
5         any     any          svc-natt  permit                          Low
                                                4
captiveportal
-------------
Priority  Source  Destination  Service          Action        TimeRange  Log  Expired
Queue  TOS  8021P  Blacklist  Mirror  DisScan  ClassifyMedia  IPv4/6
--------  ------  -----------  -------          ------        ---------  ---  -------
-----  ---  -----  ---------  ------  -------  -------------  ------
1         user    controller   svc-https        dst-nat 8081
Low                                                          4
2         user    any          svc-http         dst-nat 8080
Low                                                          4
3         user    any          svc-https        dst-nat 8081
Low                                                          4
4         user    any          svc-http-proxy1  dst-nat 8088
Low                                                          4
5         user    any          svc-http-proxy2  dst-nat 8088
Low                                                          4
6         user    any          svc-http-proxy3  dst-nat 8088
Low                                                          4
```

```
Expired Policies (due to time constraints) = 0
```

Verifying DNS Resolved IP Addresses for Whitelisted URLs

The following CLI command verifies the DNS resolved IP addresses for the whitelisted URLs in the managed device:

```
(host) #show firewall dns-names ap-name <AP-name>
Example:
(host)[md] #show firewall dns-names ap-name ap135

Firewall DNS names
------------------
Index          Name                  Id           Num-IP       List
-----          ----                  --           ------       ----
0              bugzilla              10                1    0.0.0.0
1              cricinfo               9                0
2              yahoo                  1                0
3              mycorp                 6                1    1.1.1.1
```

Bypassing Captive Portal Landing Page

An increasing number of user sessions in Captive Portal pre-authenticated role, repeatedly request the Captive Portal login page from the managed devices. This impacts the number of browser-based user login requests handled per second by the managed devices. This eventually delays the loading of the Captive Portal page and logging into Captive Portal. Most of the increased activities are from non-browser based applications running on smart phones and tablets.

Bypassing Captive Portal Landing Page is disabled by default, hence the managed devices send 200 OK status code message to the non-browser based apps.

The following CLI commands enable **Bypassing Captive Portal Landing Page** from the managed devices. When doing so, non-browser apps continue to request Captive Portal login page from the managed devices and they are responded with **302 Temporarily Moved** status code. This increases the load of the **httpd** process of the managed devices.

```
(host) [md] (config) #web-server profile
(host) [md] (Web Server Configuration) #bypass-cp-landing-page
```



The landing page contains the meta-refresh tag to reload the page using real browser applications.

Captive Portal Authentication in Bridge Mode

Starting from ArubaOS 8.7.0.0, captive portal authentication is supported for VAPs in the bridge forwarding mode. This feature supports only external captive portal servers which generate XML API/Radius CoA to the controller Only the following parameters of the **aaa authentication captive-portal** command will be supported for APs in the bridge forwarding mode:

- ap-mac-in-redirection-url
- ip-addr-in-redirection-url
- login-page
- switchip-in-redirection-url
- url-hash-key
- user-vlan-in-redirection-url

The **login-page** should be configured with an absolute path, starting with http:// or https://. This feature is supported for wireless users on all Campus AP and Remote AP models in cluster and non-cluster topology. To support captive portal authentication in the bridge forwarding mode, it is required to enable the **ageout-bridge-user** parameter in the **aaa profile** command.

Cluster is a combination of multiple managed devices working together to provide high availability to all the clients and ensure service continuity when a failover occurs.

The APs are managed by a single managed device. The client load is shared by all the managed devices. The goal of a cluster is to provide full redundancy to APs and wireless clients alike in case of a malfunction of one or more of its cluster members.

All the members in a cluster are active managed devices.

Cluster facilitates a large roaming domain, minimizes fault-domain, and helps in speedy recovery.

The objectives of a cluster are:

- **Seamless Campus Roaming:** When a client roams between APs of different managed devices within a large L2 domain, the client retains the same subnet and IP address to ensure seamless roaming. The clients remain anchored to a single managed device in a cluster throughout their roaming area which makes their roaming experience seamless because their L2 or L3 information and sessions remain on the same managed device.
- **Hitless Client Failover:** When a managed device fails, all the users fail over to their standby managed device seamlessly without any disruption to their wireless connectivity or existing high-value sessions.
- **Client and AP Load Balancing:** When there is excessive workload among the managed devices, the client and AP load is evenly balanced among the cluster members. Both clients and APs are load balanced seamlessly.

Following sections describe the pre-requisites, key considerations, and features supported in a cluster.

Requirements

Cluster is supported only on the Mobility Master and cluster members can only be managed devices.

The following managed devices support clustering:

- 7200 Series controllers - Support for up to 12 nodes in a cluster.
- 7000 Series controllers - Support for a maximum of 4 nodes in a cluster.
- 9004 controllers - Support for a maximum of 4 nodes in a cluster.
- 9012 controllers - Support for a maximum of 4 nodes in a cluster
- Mobility Controller Virtual Appliance - Support for a maximum of 4 nodes in a cluster.

Even with a 12-node cluster, the maximum supported APs and client counts are limited to 10K and 100K, respectively.

Key Consideration

Some of the key considerations are:

- All the managed devices within the cluster need to run the same software version.
- If HA-AP fast failover is enabled, then cluster cannot be enabled.
- A 12-node cluster is supported for Remote APs. Starting from ArubaOS 8.6.0.0, Remote APs can now terminate on the cluster with more than 4 nodes.
- A mix of hardware devices and the Mobility Controller Virtual Appliance-based controller is not supported.
- A Mobility Controller Virtual Appliance cluster can be set up only with same SKU models. Only homogenous clusters are supported for Mobility Controller Virtual Appliance.
- A mix of 7200 Series controllers and 7000 Series controllers within the same cluster is not recommended due to disparity in capacity between the two controller series models. However, you can use these devices in the same cluster when you want to migrate from a smaller cluster like 7000 series controllers to a larger cluster with 7200 Series controllers.
- Only homogenous cluster is supported for 9004 managed devices.
- In a cluster, the managed devices do not have to be identical.
- A managed device can be either L2- or L3-connected or it can also be a mix of both.
- Cluster is not supported for PSK-RAPs.
- Cluster is supported for external whitelist database for Remote APs in a ClearPass Policy Manager server.
- No license is required to enable the cluster feature.
- Cluster is not supported in stand-alone controllers.
- Campus APs, Remote APs, and Mesh APs are supported.
- Captive portal is not supported for the split-tunnel mode Virtual APs and wired APs, when cluster is enabled.
- A cluster supports both IPv4 and IPv6 APs in a dual-stack deployment. which is applicable to both Campus APs and Remote APs.
- Cluster is not supported when the controller-ip is set to loopback interface IP address.

Support for Homogeneous Cluster

A homogeneous cluster is a cluster built with all nodes of the same platform type.

Cluster AP Capacity

The cluster sizing depends on the number of cluster AP count required to ensure that every AP has an AAC and S-AAC with adequate capacity for all APs to failover. The recommended AP load of this cluster should be half of the total cluster capacity. Therefore, the cluster AP count should be equal to 50% of the cluster capacity.

For example, if a cluster is made up of four 7220 managed devices, the combined capacity of four 7220 managed devices is 4096 APs, hence, the AP count would be 2048.

Support for Heterogeneous Cluster

The following list provides the points to be considered for cluster capacity (APs and clients) when the cluster has a heterogeneous managed device mix. For example, 7210, 7220, and 7240 controllers.

- Total capacity of individual managed devices in the cluster, when redundancy is disabled.
- The number of cluster nodes is restricted to four when it involves a 7000 Series managed device.

- When 7200 Series managed devices are added to a cluster consisting of other 7000 Series managed devices, then the capacity of the 7200 Series managed devices is reduced to the maximum capacity of the 7000 Series managed devices that are currently part of the cluster.
- When 7000 Series managed devices are added to a cluster consisting of 7200 Series managed devices, then one of the following conditions apply:
  - If there are more than three 7200 Series managed devices in the cluster, the 7000 Series managed devices are not allowed to join the cluster.
  - If the current AP or station count on the 7200 Series managed devices is greater than the maximum AP or station capacity supported on the newly added 7000 Series managed devices, then the 7000 Series managed devices are not allowed to join the cluster. To check if the 7000 Series managed devices are allowed to join the cluster, execute the **show lc-cluster group-membership** command.
  - If the current AP or station count on the 7200 Series managed devices is lesser than the maximum AP or station capacity supported on the newly added 7000 Series managed devices, then the capacity of the 7200 Series managed devices in the cluster drops to the maximum capacity supported on the 7000 Series managed devices and the existing supported APs in the 7200 Series managed devices are not impacted.

Cluster AP Capacity

Cluster AP size should be equal to the lowest value of either 50% of total cluster capacity or the worst case scenario load. The worst case scenario load is the AP load handled by the remaining nodes in a cluster in the event of highest capacity cluster member going down.

Following examples elaborate how to calculate the cluster AP size based on the capacity of the managed devices:

Example 1:

In a cluster with one 7220 managed device and two 7240 managed devices. Capacity of a 7220 managed device is 1024 and the capacity of 7240 managed device is 2048. Now, let's calculate 50% of total capacity is (1024+2048+2048)/2 = 2560 APs. Now, assume one 7240 managed device is down, hence, the worst case scenario load is (1024 + 2048) = 3072.

Therefore, the cluster AP size in this example is 2560 APs as it is the lowest value between the 50% of total cluster capacity and the worst case scenario load.

Example 2:

In a cluster with two 7210 managed devices and one 7240 managed device. The capacity of 7210 managed device is 512 APs and the capacity of 7240 managed device is 2048 APs. So, 50% of total capacity is (512+512+2048)/2=1536 APs. Now, assume the 7240 managed device is down, hence, the worst case scenario load is (512+512) = 1024 APs.

Therefore, the cluster AP size in this example is 1024 APs as it is the lowest value between the 50% of total cluster capacity and the worst case scenario load.

Cluster Connection Types

Clustering supports the following two connection types for cluster members:

- **L2-connected**: The cluster members share the same user VLANs. All user VLANs on each node are also present in all nodes.

- **L3-connected**: The cluster members do not necessarily share the same user VLAN. Some user VLANs are not present on the other nodes.



Cluster can be formed over an L2 or L3 network. L2 is recommended for simplicity.

Roles

This section explains the roles of the members within the cluster:

Cluster Leader

When several managed devices form a cluster, the devices exchange handshake or hello messages with one another to form a cluster. When all the cluster members are in a fully connected mesh, a cluster leader is elected. The cluster leader is elected based on the highest effective priority derived from configured priority, platform value, and the MAC address of the device.

The cluster leader computes which client is mapped to which cluster member.

The cluster leader also dynamically and seamlessly balances the client load when load increases and there is an imbalance of load among the cluster members.

The cluster leader identifies standby managed devices for clients and APs to ensure hitless failover.

AAC - AP Anchor Controller

This role is given to a managed device from individual AP perspective. This is an anchor for APs. AP sets up active tunnels with its LMS-IP and also, the AAC is responsible for handling all management functions of an AP and its radios.

UAC - User Anchor Controller

This is an anchor for users. The user associates to an AP and the AP creates a dynamic tunnel to the client UAC. The UAC handles all the wireless client traffic, including association or disassociation notification, authentication, and all the unicast traffic between the managed device and the client. The UAC is used to ensure that the managed device remains the same within the cluster when clients roam between APs.

S-AAC - Standby AP Anchor Controller

A standby AAC is dynamically assigned from other cluster members. An AP sets up standby tunnels with the S-AAC. If the AAC fails, the S-AAC detects the failure and ensures that the AP fails over to the S-AAC. Dynamically, the cluster leader chooses the new S-AAC for an AP after the original AAC failed and the S-AAC becomes the new AAC.

S-UAC - Standby User Anchor Controller

This is the standby managed device from the user perspective. A user fails over to this managed device when the active UAC is down. The S-UAC is the role given to the managed device if a user fails over to this managed device when the Active UAC (A-UAC) is down.

Anchored to a Single Managed Device

A user is mapped to a UAC through a hashing algorithm at the AP level. At the AP, there is a single hashing algorithm that creates an index based on the MAC address of the client. This index points to a mapping table to the actual UAC for that user. This mapping is sent to all the nodes in the cluster by the

cluster leader and then, the AAC sends this mapping to the respective APs. So, all APs in the cluster have the same mapping information. The cluster leader assigns the S-AAC to each AP after considering the AP load on the cluster.

Remote AP Support

With Remote APs, a tunnel mode VPN is configured and each AP is assigned with an inner-IP or remote-IP. The same remote-IP or inner-IP is assigned to the Remote APs on every managed device in the cluster. Starting from ArubaOS 8.0.0.0, the cluster setup supports both IPv4 and IPv6 clients and the IPv6 clients sessions are also synchronized and continued after failovers.

Starting from ArubaOS 8.7.0.0, when both inner IPv4 address and inner IPv6 address pools are configured for Remote APs, the tunnel is established based on the outer IP address of the Remote AP. If the outer IP is IPv4 address, cluster inner IPv4 address from Remote AP inner IP pool is used to form the tunnel. Similarly, if the outer IP is IPv6 address, cluster inner IPv6 address is used to form the tunnel.

The following CLI command supports IPv4 address for Remote APs in a cluster configuration:

```
(host) [mynode] (config)#lc-rap-pool <pool_name> [{pool_start_address} {pool_end_
address}]
```

The following CLI command supports IPv6 address for Remote APs in a cluster configuration:

```
(host) [mynode] (config)#lc-rap-pool-v6 <pool_name> [{pool_start_address} {pool_end_
address}]
```

ArubaOS now provides support for ClearPass Policy Manager to whitelist Remote APs in a cluster environment. For more information, see Offloading a Controller Whitelist to ClearPass Policy Manager on page 73.

IPv6 Cluster Support

Starting from ArubaOS 8.2.0.0, IPv6 cluster is supported. Managed devices must terminate on the Mobility Master through the IPv6 IPsec tunnel.

Only IPv6 APs can terminate on an IPv6 cluster and clients can be either IPv4 or IPv6 type.

The following CLI command displays IPv6 cluster information:

```
(host) #show lc-cluster group-membership
```



VRRP-IP and VRRP-VLAN are not supported with IPv6 cluster.

ArubaOS now allows both IPv4 and IPv6 APs to connect to a cluster seamlessly in a dual-stack deployment, irrespective of the cluster IP address family. In a cluster formation, both the IPv4 and IPv6 addresses are exchanged between cluster members. Hence, the cluster can send both IPv4 and IPv6 addresses in node list to APs so that the APs are able to connect to the cluster member.

The following table provides information on the supported address modes between clusters and APs:

Table 68: *Supported Address Modes*

Cluster Address Mode	Supported on IPv4 APs	Supported on IPv6 APs
Native IPv4 cluster	Yes	No

Cluster Address Mode	Supported on IPv4 APs	Supported on IPv6 APs
Native IPv6 cluster	No	Yes
Dual-stack IPv4 cluster	Yes	Yes
Dual-stack IPv6 cluster	Yes	Yes

Cluster Features

Following sections describe the features supported on a cluster:

Enhanced Multicast Proxy

A managed device acts as a multicast proxy for all the wireless clients connected to it. The subscription of the managed device to multicast stream is done through a single VLAN. Hence, only one copy of the multicast stream will be delivered to a client.



Clustering supports only IGMP proxy and MLD.

When IGMP proxy or MLD is enabled, client reports reach the UAC. The UAC then transfers the subscription information to the AAC . Both managed devices (AAC and UAC) serve as proxies for clients in the uplink multicast VLAN.

APs are anchored on the AAC and users on the UAC. When an AP boots, it establishes a tunnel with the AAC. The same tunnel is used for UAC traffic as well. When a client comes up, the AP determines its UAC and establishes a tunnel with the UAC. When the client roams from one AAC to another, PIM detects this roaming through STA (station) channel and deletes the multicast subscriptions of the client from the old AAC and adds them to the new AAC. To perform this, a cluster proxy table that stores per-client subscriptions is maintained in the UAC.

If a multicast stream is sourced from a wireless station, the managed device forwards the stream to the multicast router through the VLAN where the client is located. The downstream is still from the multicast router to each managed device in the cluster through the configured VLAN for multicast proxy operation. If the two VLANs are the same, the proxy on the UAC of the sourcing client does not receive the stream from the multicast router.



In an L3-connected cluster, when the AAC does not have the same VLAN as the UAC, the multicast traffic from the uplink does not reach the AAC. Therefore, the cluster has to be L2-connected to stream multicast traffic.

The following CLI commands configure a cluster with multicast VLAN:

```
(host) [multicast] (cluster1) #controller 10.15.128.102 mcast-vlan
<mcast_vlan>            VLAN id
```

The following CLI command displays if a cluster is configured with multicast VLAN:

```
(host) #show lc-cluster group-profile cluster1
IPv4 Cluster Members
--------------------
CONTROLLER-IP  PRIORITY  MCAST-VLAN  VRRP-IP  VRRP-VLAN
-------------  --------  ----------  -------  ---------
10.15.128.103  128       29          0.0.0.0  0
10.15.128.104  128       29          0.0.0.0  0
```

```
10.15.128.105  128       29          0.0.0.0  0
10.15.128.102  128       29          0.0.0.0  0
Redundancy:Yes
Active Client Rebalance Threshold:50%
Standby Client Rebalance Threshold:75%
Unbalance Threshold:5%
```

Client State Synchronization

Client state synchronization feature helps resolve issues regarding seamless failover, service availability, and high availability. To achieve hitless failover, the following two conditions should be met:

- Redundancy mode needs to be enabled, this is enabled by default.
- L2 connected type, that is, the cluster members must share the same VLANs.

Stateful failover is achieved through full client synchronization from the UAC to the S-UAC. For example, the station table, the user table, the L2 user state, the L3 user state, the key cache, the PMK cache, and so on get synchronized between the UAC and the S-UAC.

Users sessions are synchronized or duplicated on an S-UAC. Only high-value sessions like FTP and DPI are synchronized. But, some sessions that are considered low value like regular HTTP traffic are not synchronized.

When there is a failover, no client is deauthenticated and hence, the client seamlessly fails over to the S-UAC .



A maximum of 10 sessions per client is supported. Client state synchronization is now supported for IPv6 clients and dual stack.

In an existing cluster, when new managed devices are added and the existing managed devices have a load more than the threshold, the load balancer ensures that traffic from UACs that are overloaded are redirected to the new managed device. In this scenario, synchronization of sessions for these users is performed before the load balancer switches the users from other UACs to ensure reliability.

Starting from ArubaOS 8.6.0.0, during a UAC failure, hitless failure of high-value application traffic such as voice is supported when the client roams between BSSIDs.

Client state synchronization is useful in two different scenarios:

- When Redundancy is OFF — When redundancy mode is turned off, a standby copy is not created for an AP or the client for failover protection. As part of load balancing, prior to planned UAC switchover, sessions are synchronized to the new UAC.
- When Redundancy is ON — When redundancy mode is turned on, the system assigns the standby managed device for all APs and clients. The sessions are synchronized to the standby UAC.

Execute the following command on one of the cluster members to view the list of duplicate users that are currently connected to S-UAC.

```
(host) #show user-table standby
```

AP LACP Support

Striping LMS IP can no longer be used to stripe the traffic as each AP has GRE tunnels to more than one managed device. Therefore, starting from ArubaOS 8.2.0.0, Cluster LACP is used to stripe traffic on a per-UAC basis. That is, in a cluster setup, the clients or users on the same AP are steered to different UACs and the traffic is striped to these UACs.

When cluster is enabled, striping IP is not used even if it is a single-node cluster; the striping of traffic for the Ethernet interfaces is according to the UAC node.

For a non-cluster setup, the striping LMS IP is used in the same way as before.

For an upstream traffic, the cluster LACP load-balances these UACs across the Ethernet ports.

For a downstream traffic, because the **Source-IP** and **MAC address** of the GRE packets are different from those of AP, the AP's uplink switch spreads the traffic.

The following CLI commands configure AP LACP in a non-cluster topology:

On an uplink switch of an AP, use the following command to configure LACP between the two ethernet ports of the AP:

```
(host) [md] (config) #ap-lacp-striping-ip
(host) [md] (AP LACP LMS map information) #aplacp-enable
(host) [md] (AP LACP LMS map information) #striping-ip 10.15.127.2 lms 10.15.127.3
```

The following CLI command displays the configuration:

```
(host) #show ap-lacp-striping-ip
AP LACP LMS map information
---------------------------
Parameter            Value
---------            -----
AP LACP Striping IP  Enabled
GRE Striping IP      10.15.127.2 LMS 10.15.127.3
```



The lms-ip value in ap-system-profile will be used as a key to look up entries in ap-lacp profile.

It is recommended not to configure GRE striping IP address for stand-alone controller deployments.

Authorization Server Interaction

This feature supports CoA requests in a cluster using multiple VRRP instances. This feature ensures that the CoA request is not dropped when the UAC changes due to controller failure or client load balancing.

CoA is change of authorization, which is an extension to RADIUS attributes and capabilities. CoA request messages are sent by a RADIUS server to a NAS device for dynamically modifying the existing session authorization attributes. A CoA-Request contains the information for dynamically changing session authorizations. If NAS is able to successfully change the authorizations of the user session(s), it responds with a CoA-ACK. Otherwise, it returns a CoA-NAK to the RADIUS server.

To support this feature, multiple VRRP instances are created dynamically, with one instance per cluster node. Here, the cluster node is the master of that instance. In a cluster, the virtual IP of each VRRP instance is used as a NAS-IP when sending RADIUS requests to the RADIUS server.



The VRRP IDs for these instances are reserved and the reserved IDs range from 220 to 255.

For example, for a cluster with 5 nodes, there are five VRRP instances and five virtual IP addresses. That is, One Virtual IP address for each VRRP instance. The cluster uses the virtual IP for an instance as the NAS-IP in a RADIUS request. That is, when the cluster node sends RADIUS requests on behalf of a client that is trying to authenticate a RADIUS server, It inserts the Virtual IP as the NAS-IP in that RADIUS packet.



VRRP VLAN can be the same as that of the controller-ip. VRRP VLAN can also be different if the same VLAN is used with all of the cluster members.

To set the VRRP IP address of the A-UAC as the NAS IP, VRRP IP must be assigned for each cluster member. This assignment process automatically configures the VRRP membership for other members of the cluster, and sets the VRRP priority correctly so that the primary A-UAC owns the virtual IP when it is up.

The following procedure describes how to set the VRRP IP address and VRRP VLAN:

1. When configuring a new cluster, select the group folder under which the managed devices are located, in the **Managed Network** node hierarchy, navigate to the **Configuration > Services > Clusters** tab.
2. Click **+** in the **Clusters** table to create a new cluster profile. The **New Cluster Profile** table is displayed.
3. Enter a name for the cluster.
4. Click **+** in the **Controllers** table to add a new controller. The **Add Controller** table is displayed.
5. Enter the **VRRP IP** and the **VRRP VLAN** field values of the managed device.
6. Click **OK**.
7. Similarly, enter the **VRRP IP** and the **VRRP VLAN** values for all managed devices.



Aruba recommends you to use the same controller-ip subnet as the VRRP-VLAN.

The following CLI commands set the VRRP IP address of the A-UAC as the NAS IP:

```
(host) [MD-cluster1]#lc-cluster group-profile primary-cluster
(host) [MD-cluster1](Classic Controller Cluster Profile "primary-cluster") #controller
10.15.43.2 vrrp-ip 100.1.1.2 vrrp-vlan 100
```

Following is an example of how to set the VRRP IP for a cluster with two managed devices:

```
(host) [MD]#lc-cluster group-profile primary-cluster
(host) [MD-cluster1](Classic Controller Cluster Profile "primary-cluster") #controller
10.15.43.2 vrrp-ip 100.1.1.2 vrrp-vlan 100
(host) [MD-cluster4](Classic Controller Cluster Profile "primary-cluster") #controller
10.15.43.5 vrrp-ip 100.1.1.5 vrrp-vlan 100
```

The following CLI commands verify the VRRP status for both managed devices:

```
(host) [MD-cluster1] #show vrrp
Virtual Router 220:
Description
Admin State UP, VR State MASTER
IP Address 100.1.1.2, MAC Address 00:00:5e:00:01:dc, vlan 100
Priority 255, Advertisement 1 sec, Preemption Enable Delay 0
Auth type NONE ********
tracking is not enabled
```

```
(host) [MD-cluster4] #show vrrp
Virtual Router 220:
Description
Admin State UP, VR State BACKUP
IP Address 100.1.1.2, MAC Address 00:00:5e:00:01:dc, vlan 100
Priority 235, Advertisement 1 sec, Preemption Enable Delay 0
```

```
Auth type NONE ********
tracking is not enabled
```

AP Failover to Different Cluster

Starting from ArubaOS 8.0.0.0, an AP can fail over between clusters. Redundancy across geographically separated data centers are supported. An AP terminates on an AAC in a cluster. If a member in the cluster fails, the AP will fails over to the S-AAC in the same cluster. If the AP is unable to establish communication with any of the members in the first cluster, then it terminates on another cluster setup in the backup data center. It terminates on another cluster only if the other cluster member IP is provided in the AP system profile as backup LMS.

For example, a cluster with four managed devices is deployed in the West Coast data center. Similarly, a cluster with four managed devices is deployed in the East Coast data center. An AP is configured to have a primary termination on the West Coast data center and backup termination on the East Coast data center. If a managed device fails in the West Coast data center, then the AAC moves to another managed device in the same data center. However, if the entire West Coast data center is inaccessible to the AP, then it fails over to the East Coast data center.

Grouping Managed Devices Within a Cluster

Starting from ArubaOS 8.2.0.0, you can group managed devices within a cluster, which helps influence the S-AAC and S-UAC assignments. The preference for both S-AAC and S-UAC is given to the managed devices in different groups compared to the group which has the AAC and UAC configured.

A new parameter, **group**, is introduced in the **lc-cluster group-profile** command.

```
(host) #lc-cluster group-profile <profile>
controller <ip> [priority <prio>] [mcast-vlan <mcast_vlan>] [vrrp-ip <vrrp_ip> vrrp-
vlan <vrrp_vlan> group <group number>]
```

AP Node List

When an AP joins a cluster, it learns the IP addresses of all the cluster members. These IP addresses are stored in a Node List, which is saved as an environment variable in the AP's flash memory. Therefore, when the AP reboots and comes back up, the AP checks the Node List, contact the cluster member that is listed first in the Node List. If the cluster member that is first on the Node List is down or not reachable, then the AP dynamically tries the second cluster member listed in the Node List and so forth. The AP always finds a managed device as long as at least one managed device is active in the cluster.



The AP rebootstraps if the entire Node List is not reachable.

APmove

This feature allows an end user to move a specific AP from the current managed device to a target managed device. The **apmove** command reassigns an AP or AP group to any managed device.

Use the **apmove** command to move a specific AP to a specific assigned managed device in the following scenarios:

- To move some specific APs to other managed device without changing any configuration.
- If there is no failover or rebootstrap configuration between the current managed device and the target managed device.

You can execute the **apmove** command in the following setups:

- Same cluster group — **apmove** can only be executed on a cluster managed device leader.
- Same HA — this command is executed on the HA-Active node and the AP fails over to HA standby.
- Normal topology — In a non-cluster setup, **apmove** can be executed on the node to move an AP from the current managed device to another managed device.

The following CLI command moves a specific AP:

If cluster is enabled, the system access point monitor process checks whether the current node is the cluster leader. If not, it displays an error and the cluster leader's IP address is provided to the end-user. The end-user can then locate the cluster leader and execute the command in the correct managed device.

The **apmove** command is executed as follows:

```
(host) [mynode] (config) #apmove <ap-mac>  <target-ip>
(host) [mynode] (config) #apmove <ap-group/all> <source-ip> <target-ip>
```

Parameter	Description
`ap-mac`	MAC address of a specific AP.
`ap-group/all`	APs in specific group or all APs in the specific managed device.
`source-ip`	Specific managed device from which the specific APs are to be moved.
`target-ip`	Specific managed device to which the APs are to be moved.

When the target IP is within the cluster, the APmove is initiated from the cluster leader. When the target IP is outside the cluster, Apmove is initiated on the AAC or S-AAC.

When APmove is initiated from the AAC, the AP gets the target IP and sets the APmove master variables. If the APmove target is a managed device outside the current cluster, then the AP rebootstraps and connects to that target managed device. Irrespective of whether the target node is in another cluster or not, the AP nodelist is purged if target IP is outside the cluster. If the target managed device is part of another cluster, then a new nodelist is sent to the AP. If the AP is unable to connect to any of the nodes in the nodelist, it falls back to other known entities such as previous_lms, backup_lms, master, and so on.

In a cluster environment, the priority given by the AP when APmove is initiated is as follows:

1. APmove master (only used in cluster upgrade scenario)
2. Cluster nodelist
3. Previous LMS (CPsec-enabled only)
4. Master variables



A nodelist is introduced to avoid multiple redirections to the AP and allows the AP to directly connect to the previous known AAC. However, if the previous known AAC is down, the AP connects to any of the nodes in the nodelist.

EST Support for Cluster

In a cluster setup, the APs establish IPsec tunnel with AAC, S-AAC, and UAC. Starting from ArubaOS 8.4.0.0, the cluster members use enrolled certificate for IPsec tunnel authentication instead of using factory certificates.

When Enrollment over Secure Transport (EST) is enabled in a cluster setup, AAC sends the EST parameters to APs and APs will undergo enrollment and establish an IPsec tunnel with all the cluster members using these enrolled certificates.

The existing cluster gets disconnected on EST activation and all the APs reboot as part of EST enrollment. During this process, the IPsec tunnels on the cluster peer are deleted, which results in the cluster getting disconnected on that peer. This ensures that the cluster traffic does not go to the peers without getting encrypted or encapsulated.



It is recommended to enable EST on all the cluster members before enabling cluster group-membership.

Configuring EST support for cluster

To configure EST support for cluster, refer to Certificate Enrollment Using EST section.

Remote AP Support with Cluster behind NAT

Remote APs were supported only with public IP addresses for all the managed devices in a cluster deployment. But, the cluster behind NAT cannot work with Remote APs because the managed devices in the cluster use switch IPs which are in private domain; to which the Remote AP does not have access.

Starting from ArubaOS 8.4.0.0, Remote APs can map the managed device's private address to a public space by obtaining the private IP and public IP address mapping from a cluster. Therefore, the cluster behind NAT is supported with Remote APs.

Key Consideration

- Remote APs are provisioned with any of the public IP address that the cluster is using.
- NAT mapping is configured in the customer NAT device accordingly to what the cluster profile is using
- The mapping must be allowed even if a firewall is configured.

Limitations

- Configuration of same public IP for different nodes in the same cluster profile is not allowed.
- Configuring same public IP across different cluster profiles only when one profile is active across all cluster members.
- Cluster is not supported for external whitelilst-db.



Mapping between the public and private addresses configured in the cluster profile must be configured in the NAT device as well.

The following procedure describes how to enable a Cluster behind NAT with Remote APs:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > Clusters** tab.
2. Click **+** in the **Clusters** table to create a new cluster profile. The **New Cluster Profile** table is displayed.
3. Enter the cluster name, **rapcluster**.
4. Enter the RAP Public IP along with the parameters listed in
5. Click **Submit**.
6. Then, in the **Cluster Profile** tab, select **rapcluster** from the **cluster group-membership** drop-down list.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes.**

The following CLI commands map the public and private addresses with the Remote AP in a cluster profile:

```
(host) [cluster] (config) #lc-cluster group-profile  rapcluster
(host) [cluster] (Classic Controller Cluster Profile "rapcluster") controller
10.10.10.1 rap-public-ip 100.100.100.101
(host) [cluster] (Classic Controller Cluster Profile "rapcluster")controller
10.10.10.2 rap-public-ip 100.100.100.102
(host) [cluster] (Classic Controller Cluster Profile "rapcluster")controller
10.10.10.3 rap-public-ip 100.100.100.103
(host) [cluster] (Classic Controller Cluster Profile "rapcluster")controller
10.10.10.4 rap-public-ip 100.100.100.104
```



When this profile is configured in the group-membership, then the corresponding public IP for that cluster member is used.

The following CLI commands check if the public IP of the Remote AP is configured based on the controller's private IP address:

```
(host) #Show lc-cluster group-profile

IPv4 Cluster Members
--------------------
CONTROLLER-IP  PRIORITY  MCAST-VLAN  VRRP-IP  VRRP-VLAN  GROUP-ID  RAP-PUBLIC-IP
-------------  --------  ----------  -------  ---------  --------  -------------
10.17.62.194   128       0           1.1.1.1  200        0         10.10.10.11
10.17.62.195   128       0           1.1.1.2  200        0         10.10.10.12
```

VRRP ID and Passphrase

Cluster allows users to set the starting value of VRRP ID and passphrase for a virtual IP in the cluster profile to avoid VRRP conflict. That is, Cluster VRRP members will be assigned consecutive VRRP IDs starting from the value configured.

Traditionally, when a user configured a virtual IP in a cluster, ArubaOS automatically configured the VRRP groups between the range, 220 - 225. This lead to VRRP conflicts when multiple clusters shared the

same L2 network. Therefore, to avoid VRRP conflict, clusters now allow users to set the VRRP ID for a virtual IP in the cluster profile.

Following parameters can be set by the user in the cluster configuration profile:

- Specify the starting VRRP ID
- Specify the VRRP passphrase for securing the VRRP session

The following CLI commands configure the VRRP ID and VRRP passphrase:

```
lc-cluster group-profile <profile-name>
vrrp-id <starting id>  [   vrrp-passphrase <vrrp passphrase string>]
```

Parameter	Description
`vrrp-id`	This is an optional parameter which specifies the starting VRRP ID for cluster members. If this is not configured, system automatically configures VRRP groups within the range of 220-225.
`vrrp-passphrase`	This is an optional password of up to 8 characters that can authenticate VRRP peers in their advertisements. If this is not configured, there is no authentication password.

The following CLI command checks the configuration:

```
(host) #show lc-cluster group-profile v4cluster
IPv4 Cluster Members
--------------------
CONTROLLER-IP  PRIORITY  MCAST-VLAN  VRRP-IP  VRRP-VLAN  GROUP-ID  RAP-PUBLIC-IP
-------------  --------  ----------  -------  ---------  --------  -------------
10.20.101.12   128       0           0.0.0.0  0          0         0.0.0.0
10.20.101.5    128       0           0.0.0.0  0          0         0.0.0.0
10.20.101.20   128       0           0.0.0.0  0          0         0.0.0.0
10.20.101.7    128       0           0.0.0.0  0          0         0.0.0.0
Redundancy:Yes
Active Client Rebalance Threshold:20%
Standby Client Rebalance Threshold:40%
Unbalance Threshold:5%
Active AP Load Balancing:YES
Active AP Rebalance Threshold:20%
Active AP Unbalanced Threshold:5%
Active AP Rebalance Count:50
Active AP Rebalance Timer:1 mins
Starting VRRP ID:99
VRRP Passphrase:********
```

Cluster Configuration

This section describes the procedure for setting up a cluster and editing a cluster profile using the WebUI and the CLI.

Configuring a Cluster

Following section describes how to configure a cluster using the WebUI. The configuration is carried out in two stages:

- Creating a cluster profile.
- Attaching the created profile to the cluster group membership.

Perform the following steps to add a cluster profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > Clusters** tab.
2. Click **+** in the **Clusters** table.
3. Enter a name for the cluster profile in the **Name** field.
4. Click **Submit**.
5. To configure the cluster created, select the cluster from **Clusters** table.
6. In the **Cluster Profile > <cluster name>** window, expand **Basic**.
7. To add controllers to the cluster, click **+** in the **Controllers** table. The **Add Controller** window is displayed.
8. Define the parameters listed in Table 69 .
9. Click **OK**.
10. Expand **Advanced**.
11. Select the **Redundancy** check box to enable redundancy in the cluster.
12. Optionally, the **Active client rebalance threshold, Standby client rebalance threshold, Unbalance threshold, and Heartbeat threshold** can be set. However, these parameters have default settings and Aruba strongly recommends you to use the default settings.



For **Minimum Heartbeat Threshold in milliseconds**, the default value for a port channel is 2000 msec and for a single Ethernet connection (without port channel) is 900 msec. However, if heartbeat threshold is configured to a custom value, then that value takes precedence over the default values.

13. Click **Submit**.
14. Perform the following steps to attach the cluster profile to the cluster group membership.
15. In the **Managed Network** node hierarchy, select a managed device that you want to add to the cluster.
16. Navigate to the **Configuration > Services > Cluster** tab and expand **Cluster profile**.
17. Select a cluster profile from the **Cluster group-membership** drop-down list.
18. Set the **Exclude VLAN** field by either typing or selecting from the drop-down list to build a list of VLAN IDs separated by commas.



In the **Exclude VLAN** drop-down list, if the user selects a VLAN ID, the selected value gets added to the already existing content in the field. For example, if the text field contains '2' and the user selects '5' from the drop-down list, the field must display '2,5'. A range of value can also be added, for example, 1-5.

19. Click **Submit**.
20. Click **Pending Changes**.
21. In the **Pending Changes** window, select the check box and click **Deploy changes.**

Table 69: *Cluster Profile Parameters*

Parameter	Description
IP version	Select the IP version - IPv4 or IPv6.
IP address	The IP address must be set to the switch IP of the managed device.
Group	This is used to influence the S-UAC and S-AAC assignments made by the cluster leader. Enter an integer value between 1 and 12 for the group id.
VRRP IP	The IP used to service all requests initiated by external authentication servers such as CoA.
VRRP VLAN	The VLAN used to service all requests initiated by the external authentication servers such as CoA
MCast VLAN	The VLAN used to subscribe the multicast traffic to the upstream multicast router.
Priority	This is used to influence the cluster leader election.

The following CLI commands set up a cluster:

1. To create a cluster node:

```
(host) [mynode] (config) #configuration node /md/cluster
```

2. To change to the configuration cluster node that you created:

```
(host) [mynode] (config) #change-config-node /md/cluster
```

3. To configure a managed device under the previously created node.



Ensure that the common profiles such as SSID, VAP , and AAA profiles configured in /managed device/cluster are consistent.

```
(host) [mynode] (config) #configuration device 00:1a:1e:02:04:88 device-model
A7210 /md/cluster
```

4. Repeat this configuration for multiple managed devices.
5. All managed devices in the cluster need to be time-synchronized. Hence, it is recommended to have an NTP server in a cluster setup. To configure an NTP server:

```
(host) [cluster] (config) #ntp server <ip address> iburst
```

```
(host) [cluster] (config) #ntp authentication-key 1 md5 <password>
```

6. To configure the cluster group profile in the Mobility Master:

```
(host) [cluster] (config) #lc-cluster group-profile 6NodeCluster
```

7. Managed devices IP addresses in lc-cluster group-profile can be either IPv4, or IPv6, or a combination of both. However, on the Mobility Master, we can configure IPv4 cluster and IPv6 cluster separately. Both clusters function independently and the Mobility Master can send the configuration updates to the respective managed device.

8. To add the managed devices to the group profile:



The switch IP of the managed device is used as the IP address in the following configuration. The AP's termination point must also be set to the switch IP of the managed device. The LMS-IP for the AP in the AP system profile becomes the active-AAC (A-AAC) for the AP.

9. For IPv6 network:

```
(host) [cluster] (Classic Controller Cluster Profile "6NodeCluster ") #vrrp-vlan 0 group 200:192:168:28::24 priority 128 mcast-vlan 0 vrrp-ip-v6 ::
(host) [cluster] (Classic Controller Cluster Profile "6NodeCluster ") #vrrp-vlan 0 group 200:192:168:28::26 priority 128 mcast-vlan 0 vrrp-ip-v6 ::
(host) [cluster] (Classic Controller Cluster Profile "6NodeCluster ") #vrrp-vlan 0 group 200:192:168:28::22 priority 128 mcast-vlan 0 vrrp-ip-v6 ::
(host) [cluster] (Classic Controller Cluster Profile "6NodeCluster ") #vrrp-vlan 0 group 200:192:168:28::23 priority 128 mcast-vlan 0 vrrp-ip-v6 ::
```

10. For IPv4 network:

```
(host) [cluster] (Classic Controller Cluster Profile "6NodeCluster ") #vlan 0 group 1 192.168.28.22 priority 128 mcast-vlan 0 vrrp-ip 0.0.0.0 vrrp-
(host) [cluster] (Classic Controller Cluster Profile "6NodeCluster ") #vlan 0 group 1 192.168.28.23 priority 128 mcast-vlan 0 vrrp-ip 0.0.0.0 vrrp-
(host) [cluster] (Classic Controller Cluster Profile "6NodeCluster ") #vlan 0 group 2 192.168.28.24 priority 128 mcast-vlan 0 vrrp-ip 0.0.0.0 vrrp-
(host) [cluster] (Classic Controller Cluster Profile "6NodeCluster ") #vlan 0 group 2 192.168.28.26 priority 128 mcast-vlan 0 vrrp-ip 0.0.0.0 vrrp-
```



IP address is a mandatory parameter and **priority**, **group**, **mcast**, **VLAN**, **VRRP IP,** and **VRRP VLAN** are optional parameters.

11. In the Mobility Master, apply the configuration to managed devices:

```
(host) [cluster] (Classic Controller Cluster Profile "6NodeCluster ") #write memory
```

12. Saving Configuration...
13. Partial configuration for /md/cluster
14. Configure the group-membership on each managed devices. If you have nodes only under a node-path that forms a cluster, then execute the command on that node-path [00:1a:1e:02:04:88].

```
(host) [00:1a:1e:02:04:88] (config) #lc-cluster group-membership 6NodeCluster
(host) [00:1a:1e:02:04:88] (config) #write memory
```

15. On each managed device, check the cluster status:

```
(host) #show lc-cluster group-membership
```

16. To ensure the correct working of client SSO upon failover, managed devices in the cluster must be L2-connected. The following command shows the status of L2 or L3 connectivity in cluster.

```
(host) [md] (cluster)#show lc-cluster vlan-probe status
```

17. Optionally, on the managed devices, exclude certain VLANs for the VLAN probing algorithm.

```
(host) (config) #lc-cluster exclude-vlan <vlan-number>
```

18. After removing the VLANs using the previous command, run the VLAN probing algorithm again.

```
(host) [cluster] (config) #lc-cluster start-vlan-probe
```

Editing a Cluster Profile

The following procedure describes how to edit a cluster profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > Clusters** tab.
2. To edit an existing managed device, select the managed device from the Controllers list. To add managed devices to the cluster, click **+** in the **Controllers** table.
3. Edit or enter the values for the parameters described in Table 69.
4. Click **OK**.
5. Expand **Advanced** to edit parameters for Active AP load balancing described in Table 70. However, these parameters have default settings and Aruba strongly recommends you to use the default settings.



When an infrastructure network is not able to handle the load, cluster heartbeat timeout can happen. To handle this, either prioritize the cluster heartbeat packets on the infrastructure network or increase the heartbeat timeout on the cluster profile.

6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Using Basic Show Commands

Use the following **show** commands to ensure that the cluster configuration is working as expected:

Check the cluster status on each managed device:

```
(host) #show lc-cluster group-membership
```

View the status of the VLAN probing algorithm, which runs automatically between every pair of nodes in cluster:

```
(host) #show lc-cluster vlan-probe status
```

View the reason for cluster member disconnection due to various events and to view the last time stamp of disconnection:

```
(host) # show lc-cluster heartbeat counters
```

View the cluster heartbeat counters:

```
(Host) #show datapath cluster heartbeat counters
```

View the history of the connection and disconnection events:

```
(host)show lc-cluster history
```

View the active or standby AP load distribution within the cluster for an AP:

```
(host) # show lc-cluster load distribution ap
```

View the active or standby client load distribution within the cluster for a client:

```
(host) # show lc-cluster load distribution client
```

View the list of APs in standby mode on managed devices:

```
(host) # show ap standby
```

View the list of users in standby mode on managed devices:

```
(host) # show user-table standby
```

View the list of users in datapath in standby mode on managed devices:

```
(host) # show datapath user standby
```

View the A-UAC and S-UAC for any given client. This command can be run on any managed device that is part of the cluster:

```
(host) # show aaa cluster essid <essid name> mac <client mac address>
```

View the detailed information about all the connected peers including the heartbeat requests sent or responses received, all the sequence number of missed and delayed heartbeats along with time stamp , last received or sent sequence number to dead peer and the time stamp, This command also displays the current cluster member's heartbeat threshold and threshold updated count, added or deleted peer count and the current time stamp.

```
(Host) #show datapath cluster details
```

Collect the cluster-related debug information from managed devices:

```
 (host) #show cluster-tech-support </flash/config/outfile>
```

Collect the cluster-related debug information from an AP:

```
(host) #show ap cluster-tech-support ap-name <ap-name> </flash/config/ap outfile>
```

Collect the IPv6-related debug information:

```
(host) #show gsm debug channel sectun
```

View the Remote AP inner IP pool for cluster deployment:

```
(host) #show lc-rap-pool rap-cluster
```

Cluster Load Balancing

Cluster load-balancing is achieved through the features, Client load-balancing and AP load-balancing. Both these features are explained in this section.

Client Load Balancing

The client load balancing feature ensures that clients are evenly distributed across the cluster members, thereby using the system resources efficiently.

If the system detects a distorted distribution of load, it balances the load on the managed devices by changing the UAC of these clients. The load across all the managed devices is balanced in the cluster regardless of the type of platform.

The cluster manager calculates the ratio of the existing number of clients on a managed device and its maximum capacity. Based on this ratio and additional threshold triggers, client load balancing is triggered.

When any new managed device, including the managed device that comes up after a failover, is added to an existing cluster, it is considered for load balancing and accordingly, APs and clients are moved to balance the load in the cluster.



Load balancing is enabled by default when a cluster is configured.

Threshold triggers

- **Active client rebalance threshold:** The actual active load on a cluster member. The threshold is set at 20%, that is, 20% of the capacity of a platform.
- **Standby client rebalance threshold:** The standby load on a cluster member. The threshold is set at 40%.
- **Unbalanced threshold:** The difference between the loads on maximum loaded cluster node and the minimum loaded cluster node. The threshold is set at 5%, that is, there must be at least a 5% disparity in load between the managed devices.
- **AP Total Load Balance threshold:** The total load balance threshold is set to 40%. This is the default value and cannot be configured.

For load balancing to be triggered for active clients, the active client rebalance threshold and the unbalanced threshold percentages must be met. Similarly, for the standby client, the standby client rebalance threshold and the unbalanced threshold percentages must be met.



When the redundancy mode is enabled, the capacity of the cluster is reduced to half.

AP Load Balancing

The AP load balancing feature ensures that the cluster leader manages the load balancing based on the platform capacity. The AP is dynamically assigned an AAC when it connects to a cluster. Here, instead of client load, AP load is considered.

Both active and standby APs are considered for load balancing.

Following is the AP load balancing criteria if a managed device is newly added:

- When an AP threshold is already met in the cluster nodes, if a new managed device is added, the Active AP table of the new managed device is filled first based on AP count set.
- When the threshold is not met, APs are moved to standby AP table of the newly added managed device.
- The count of these APs will increment based on the AP count set only after the stabilization of the cluster, however, the APs that were moved during this phase cannot be always based on AP Count.



Starting from ArubaOS 8.3.0.0, the Active AP load balancing feature is enabled by default. In previous releases, this feature is disabled by default.

The active AP load balancing is performed using VRRP for L2 connection and the switch IP of one cluster member for L3 connection. The **AP Total Load Balance threshold** is set to 40% and **active AP load balancing threshold value** is set to 20%. This is the default value.



The LMS IP is ignored when the AP establishes communication with the cluster. However, in case of a failover, the backup-LMS IP is used.

Prior to ArubaOS 8.3.0.0, the cluster leader considered the AP load on each cluster member and assigned the cluster member with least total AP load as the AAC.

For initial load balancing, cluster leader evaluates if the managed device with the least active AP load percentage can accommodate additional APs. If yes, return this managed device as the candidate AAC.

For periodic load balancing, the cluster leader performs load balancing based on the following conditions:

1. Finds the managed devices with maximum and minimum active load percentage.
2. Finds the managed devices with maximum and minimum total load percentage
3. Checks if the max active load percentage is more than the active load percentage threshold. Also, checks if the difference between the maximum and minimum loaded managed devices is more than the unbalance threshold.
4. Moves the active APs from maximum loaded managed devices to minimum loaded managed device. However, if it is unable to move the Active APs, it will re-balance the standby load by moving the standby APs from the maximum load percentage managed device to minimum load percentage managed device.



The periodic load re-balancing occurs every 1 minute, which is the default value and APs are considered for load re-balancing based on AP count, which is 50 by default.

After APs are listed on a cluster member, the cluster manager periodically recalculates the load of the cluster members to balance the load. For example, when a new managed device joins the cluster.



The triggers for client load balancing and AP load balancing are same.

Listed below are the advantages of AP load balancing

- Easy scaling of cluster nodes.
- Eliminates manual distribution through the LMS-IP.

Configuring Cluster Load Balancing

The following procedure describes how to configure load balancing for a cluster:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > Clusters** tab.
2. In the **Clusters** table, select a cluster to configure AP load balancing.
3. In the **Cluster Profile > <cluster name>** window, expand **Advanced**.
4. Configure the active AP load balancing settings described in Table 70.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes.**

Table 70: *Active AP Load Balancing Parameters*

Parameter	Description
Redundancy	Enable redundancy. When the redundancy mode is enabled, the capacity of the cluster is reduced to half and only 8000 clients are considered to reach the threshold.
Active client rebalance threshold	Indicates the minimum total load percentage required to perform AP load balancing in the cluster. At least one of the managed devices must have this value as the total load percentage to trigger AP load balancing.
Standby client rebalance threshold	Indicates the minimum total load percentage required to perform standby AP load balancing in the cluster.
Unbalance threshold	If a managed device reaches the total load threshold, AP load balancing is triggered. This happens when the difference between the loads on the maximum loaded managed device and the minimum loaded managed device in the cluster exceeds the unbalance threshold value.
Heartbeat threshold	Minimum heartbeat threshold is set in milliseconds. The default setting is based on the latency determined between each pair of managed devices and the cluster. It also depends on the connection type between managed device and distribution switch (single ethernet cable, or port channel, and so on).

The following procedure describes how to configure the AP and client load balancing values:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System** tab.
2. Click **Profiles**. Expand **Cluster** and click **Classic Controller Cluster**.
3. Click + to create a cluster profile.
4. Add the **AP load balancing** values and **Client Load Balancing** values.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes.**

The following CLI commands configure load balancing for a cluster:

Following is an example of Active AP load balancing:

```
(7210-24) #show lc-cluster load distribution ap
Cluster Load Distribution for APs
---------------------------------
Type IPv4 Address    Active APs     Standby APs
---- --------------- -------------- ---------------
peer   192.168.28.23             25              20
self   192.168.28.24             20              25

Total: Active APs 45 Standby APs 45

(host) #show lc-cluster group-membership
Cluster Enabled, Profile Name = "ap-lb"
Redundancy Mode On
Active Client Rebalance Threshold = 20%
Standby Client Rebalance Threshold = 40%
Unbalance Threshold = 5%
AP Load Balancing: Enabled
Active AP Rebalance Threshold = 20%
```

```
Active AP Unbalance Threshold = 5%
Active AP Rebalance AP Count = 50
Active AP Rebalance Timer = 1 minutes
Cluster Info Table
------------------
Type IPv4 Address    Priority Connection-Type STATUS
---- --------------- -------- --------------- ------
peer   192.168.28.23      128    L2-Connected CONNECTED (Leader, last HBT_RSP 10ms
ago, RTD = 0.000 ms)
self   192.168.28.24      128             N/A CONNECTED (Member)
```

To check if Active AP load balancing is enabled or disabled:

Example for IPv4:

```
(host) #show lc-cluster group-membership
Cluster Enabled, Profile Name = "testLB"
Redundancy Mode On
Active Client Rebalance Threshold = 20%
Standby Client Rebalance Threshold = 40%
Unbalance Threshold = 5%
AP Load Balancing: Enabled
Active AP Rebalance Threshold = 20%
Active AP Unbalance Threshold = 5%
Active AP Rebalance AP Count = 50
Active AP Rebalance Timer = 5 minutes
Cluster Info Table
------------------
Type IPv4 Address     Priority Connection-Type  STATUS
---- --------------- --------  ---------------  ------
self 192.168.10.38    128      N/A              CONNECTED (Leader)
peer 192.168.10.34    128      L2-Connected     CONNECTED (Member, last HBT_RSP 38ms
ago, RTD= 0.000 ms)
```

Example for IPv6:

```
(host) #show lc-cluster group-membership
Cluster Enabled, Profile Name = "72xx"
Redundancy Mode On
Active Client Rebalance Threshold = 20%
Standby Client Rebalance Threshold = 40%
Unbalance Threshold = 5%
AP Load Balancing: Enabled
Active AP Rebalance Threshold = 20%
Active AP Unbalance Threshold = 40%
Active AP Rebalance AP Count = 50
Active AP Rebalance Timer = 1 minutes
Cluster Info Table
------------------
Type IPv6 Address        Priority   Connection-Type   STATUS
---- ------------------- ---------   ---------------  -------------------------
peer 2000:192:168:28::24 128        L2-Connected      CONNECTED (Member,last HBT_RSP
68ms ago, RTD = 0.000 ms)
peer 2000:192:168:28::26 128        L2-Connected      CONNECTED (Member,last HBT_RSP
66ms ago, RTD = 0.000 ms)
peer 2000:192:168:28::22 128        L2-Connected      CONNECTED (Member,last HBT_RSP
69ms ago, RTD = 0.503 ms)
self 2000:192:168:28::23 128        N/A CONNECTED    (Leader)
```

To display the number of times a managed device published APs on the AP channel for A-AAC assignment:

```
(host) #show ap debug gsm-counters
```

```
STM GSM Counters
----------------
Name                                                    Value
----                                                    -----
AP Publish Events                                       93
AP Delete Events                                        29
AP Publish Events(Load Balance)                         2
AP Delete Events(Load Balance)                          2
Radio Publish, Activate, Activate Errors                28 11 0
Radio Delete Events                                     3
Radio Delete Errors                                     15
BSS Publish Events                                      41
Responses to BSS Rcvd                                   41
BSS Delete Events                                       18
BSS Delete Errors                                       30
BSS Delete Key Not Found (included above)               30
STA Publish Events                                      0
STA Delete Events                                       0
STA Activate on S-UAC, Errors                           0 0
STA Activate for Delete, Errors                         0 0
WIRED_AP Publish Events                                 0
```

To display the number of times the AP tried to connect to the cluster leader and needed the cluster leader to assign an A-AAC for this AP:

```
(host) #show lc-cluster gsm counters | exclude 0
Cluster GSM Channel Counters
----------------------------
AP Channel: Adds                                                >> 1
AP Channel: Deletes                                             >> 1
BSS Channel: Adds                                               >> 2
BSS Channel: Section Update                                     >> 2
AP Channel: Adds and Need AAC Assignment                        >> 1
AP Channel: Deletes from SAPM, AP redirected                    >> 1
```

Cluster Deployment Scenarios

Clusters can be deployed in four different scenarios. The following section describes the guidelines for these different cluster deployment scenarios.

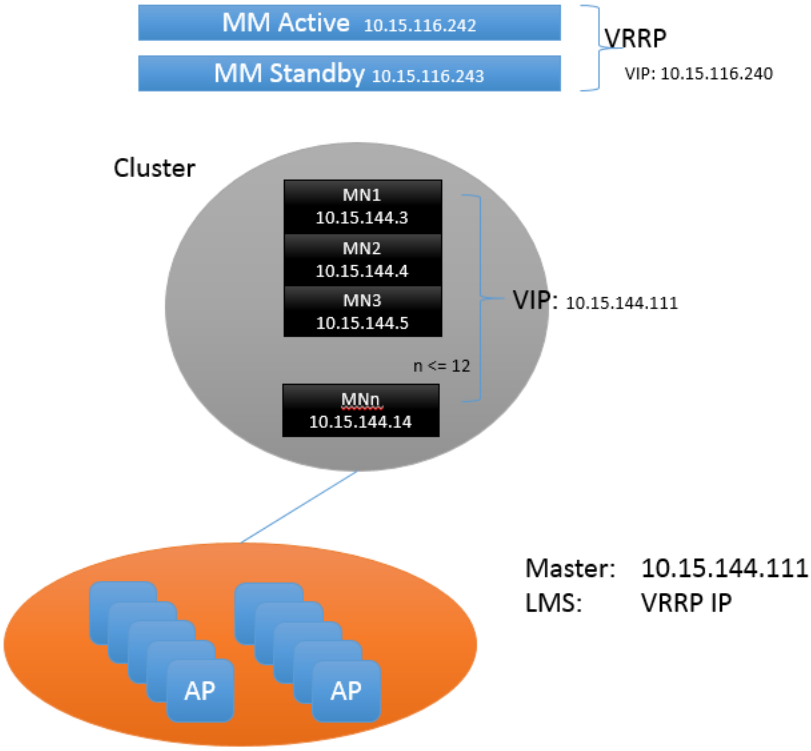


Aruba recommends you to enable AP load balancing for cluster using the CLI command **active-ap-lb**. It is enabled by default.

If **active-ap-lb** is disabled, LMS is used for the initial termination. LMS preemption is used as described in all scenarios mentioned.

Scenario 1: Cluster with Virtual IP Setup

In this scenario, an AP performs a cluster failover to the S-AAC if the A-AAC (LMS) is down. The APs perform internal rebootstrap if both A-AAC and S-AAC are down at the same time. If the AP reboots on any node including the LMS, the AP remembers the nodelist and tries all the entries in the nodelist. The AP performs a legacy rebootstrap only when it cannot reach any of the nodes.

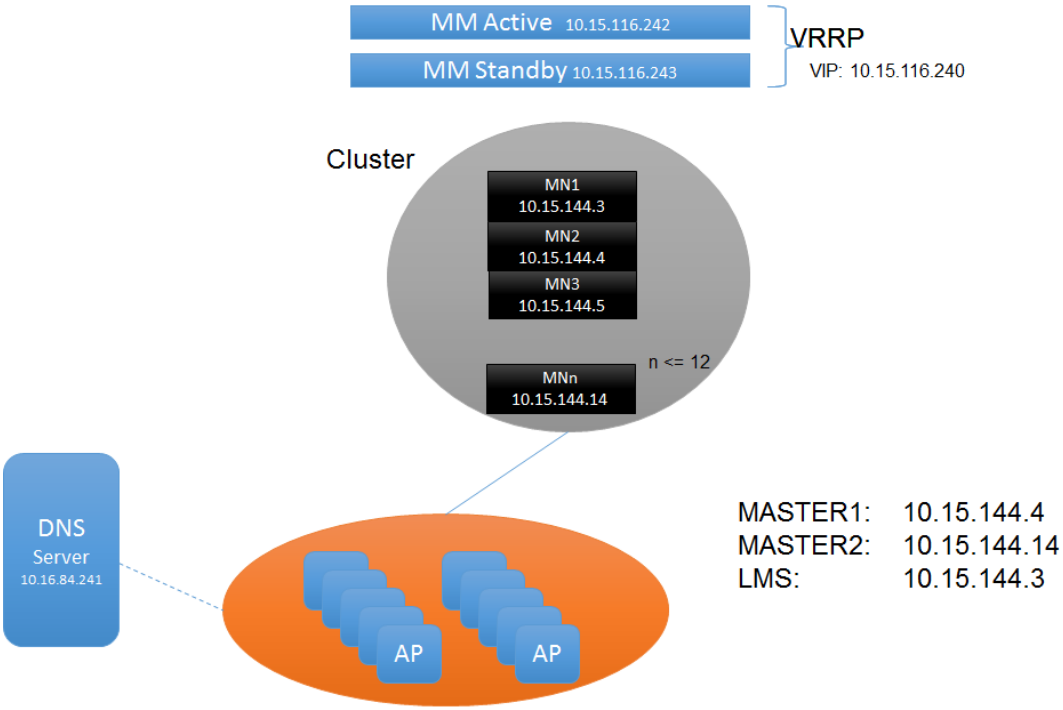


Following are the guidelines to ensure a successful deployment of the cluster in a Virtual IP :

- Master of the APs must be configured as the virtual IP on the cluster nodes.
- If the cluster has VRRP IP configured, set the VRRP IP in LMS IP address and backup-LMS IP addresses.
- Nodelist from the cluster node is saved on the AP. If the A-AAC and S-AAC are down at the same time, the AP will perform an internal rebootstrap and tries different nodes from the nodelist till the nodelist is exhausted.

Scenario 2: Cluster with Multiple Master via DNS resolution

In this scenario, an AP will perform a cluster failover to the S-AAC if the A-AAC (LMS) is down. The AP internally rebootstraps if both A-AAC and S-AAC are down at the same time and the AP tries to contact another node in the cluster till it is unable to reach the entire nodelist in the cluster. If the AP reboots on any node including the LMS, the AP remembers the nodelist and tries all the entries in the nodelist. The AP performs a legacy rebootstrap only when it cannot reach any of the nodes.



Following are the guidelines to ensure for the successful deployment of the cluster in a multiple master via DNS resolution setup:

- APs must get multiple masters using the DNS resolution.
- If the cluster has VRRP IP configured, set the VRRP IP in LMS IP address and backup-LMS IP addresses.
- Nodelist from the cluster node is saved on the AP. If both A-AAC and S-AAC are down at the same time, the AP performs an internal rebootstrap and tries different nodes from the nodelist till the nodelist is exhausted.

NOTE

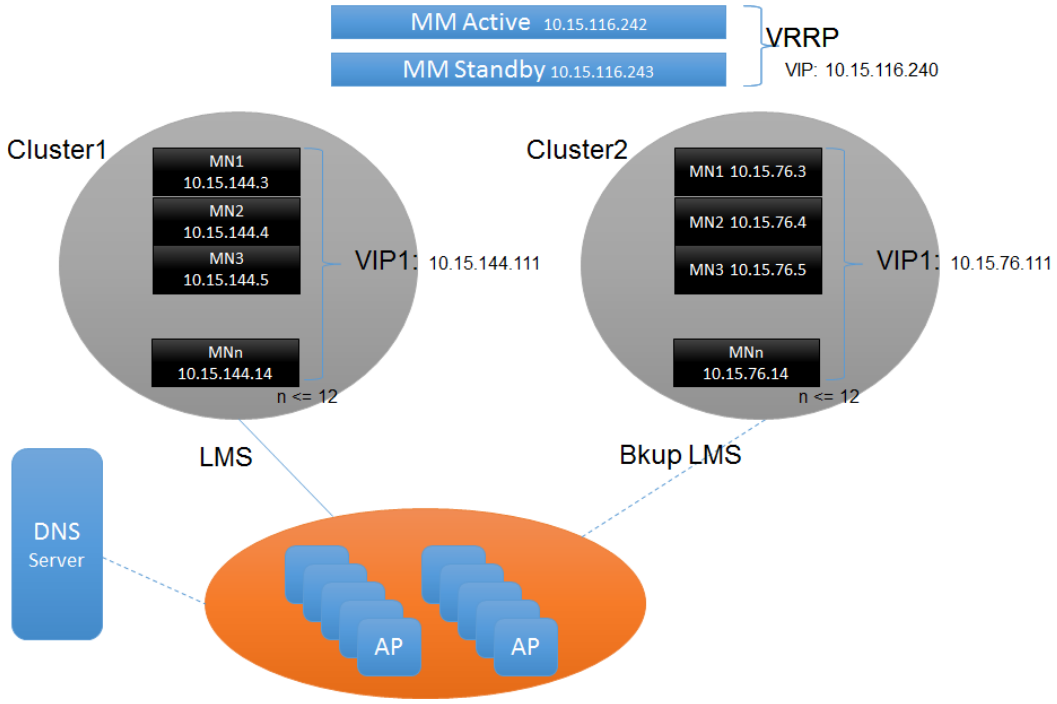
In a large deployment, Aruba recommends this configuration to avoid large failure domain.

Scenario 3: Cluster with Virtual IP via DNS Resolution Across Data Centers

In this scenario, when an A-AAC is down, the AP fails over to an S-AAC. An AP internally rebootstraps if both A-AAC and S-AAC are down at the same time and the AP tries to contact another node in Cluster1 till all the nodes are exhausted in the Cluster1 nodelist. If the AP is not able to reach Cluster1, it fails over to the backup LMS.

If LMS preemption is enabled, APs preempt to Cluster1 when the primary LMS node is up on Cluster1. The APs remain on Cluster2 if the LMS preemption is disabled even though the Cluster1 is up.

If the AP reboots on any node including the LMS, the AP remembers the nodelist and tries all the entries in the nodelist. The AP performs a legacy rebootstrap only when it cannot reach any of the nodes.



Following are the guidelines to ensure a successful deployment of the cluster with Virtual IP via DNS resolution across data centers:

- AP boots up and has two masters (one from each cluster) resolved from the DNS server. The master of the AP is resolved to virtual IP of Cluster1 and virtual IP of Cluster2.
- Nodelist from the cluster node is saved on the AP. If both A-AAC and S-AAC are down at the same time, the AP performs an internal rebootstrap and tries different nodes from the nodelist till the nodelist is exhausted.
- If AP load balancing is disabled, LMS of the ap-group or ap-name must be the IP address of the Cluster1 node and the backup-LMS must be the IP address of the other node in the Cluster2. That is, LMS of the ap-group or ap-name must be configured to the Cluster1 node and the backup-LMS must be configured to the Cluster2 node.

NOTE

In a large deployment, Aruba recommends this configuration to avoid large failure domain.

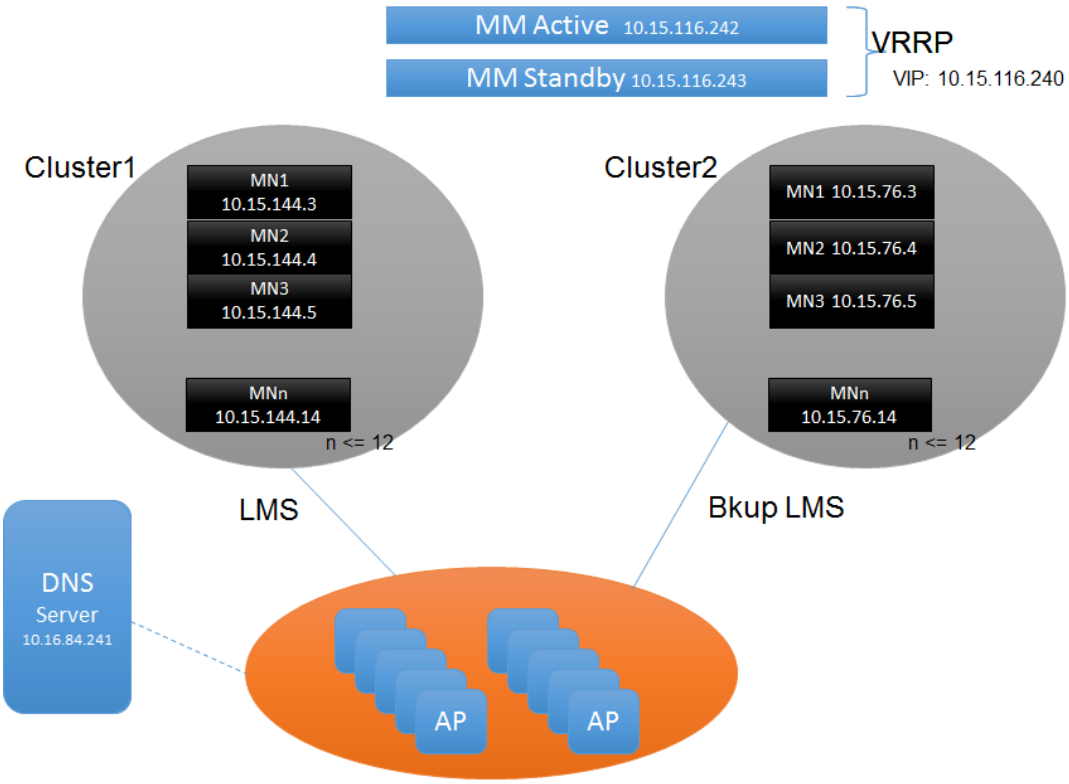
Scenario 4: Cluster with Multiple Master via DNS Resolution Across Data Centers

In this scenario, when an A-AAC is down, the AP fails over to an S-AAC. An AP internally rebootstraps if both A-AAC and S-AAC are down at the same time and the AP tries to contact another node in Cluster1 till all nodes in the nodelist of Cluster1 are exhausted. If the AP is unable to reach Cluster1, it fails over to the backup LMS.

APs terminate on the node of Cluster2, which is configured as a backup-LMS using legacy failover.

If LMS preemption is enabled, APs will preempt to Cluster1 when the primary LMS node is up on Cluster1. APs remain on Cluster2 if the LMS preemption is disabled even though the Cluster1 is up.

If the AP reboots on any node including the LMS, the AP remembers the nodelist and tries all the entries in the nodelist. The AP will perform a legacy rebootstrap only when it cannot reach any of the nodes.



Following are the guidelines to ensure a successful deployment of the cluster with multiple master via DNS resolution across data centers:

- AP boots up and has four masters (two from each cluster) resolved from the DNS server. The master of the AP must be resolved to two nodes in Cluster1 and two nodes in Cluster2.
- Nodelist from the cluster node is saved on the AP. If both A-AAC and S-AAC are down at the same time, the AP performs an internal rebootstrap and will try different nodes from the nodelist till the nodelist is exhausted.
- If AP load balancing is disabled, LMS of the ap-group or ap-name must be the IP address of the Cluster1 node and the backup-LMS should be the IP address of the other node in the Cluster2. That is, LMS of the ap-group or ap-name must be configured to the Cluster1 node and the backup LMS must be configured to the Cluster2 node.

NOTE

Maximum of 10 entries are supported for master resolution. Combination of nodes from Cluster1 and Cluster2 can be used in DNS master resolution.

Upgrading Cluster

The Live Upgrades feature allows you to upgrade the managed devices and APs in a cluster to the latest ArubaOS version. Managed devices in a cluster can be seamlessly upgraded by specifying the new image file and a target partition. This is a real-time network upgrade where managed devices and APs upgrade automatically without any planned maintenance downtime.

You can also schedule an upgrade to a specified time to avoid manual intervention. The cluster is upgraded automatically at the scheduled time. You can view, delete, or reschedule the scheduled cluster upgrade.

The live upgrade process is optimized to take less time to upgrade all the devices. The upgrade process skips mesh point and mesh portal upgrade during live upgrade but these APs will go standard upgrade when the last managed device in the cluster successfully upgrades.

When a cluster is upgraded, the following actions take place:

1. The cluster upgrade manager sends information of the APs to AirMatch.
2. The AirMatch creates logical groups of APs and updates the cluster upgrade manager with this information.
3. A target managed device is assigned to all APs in each partition.
4. The managed devices download the new firmware. The following events happen after the firmware download:
   a. For each managed device, irrespective of the partitions, AP image preload is initiated for a small capacity of the AP platform. For example, 1/8th capacity of the AP platform at a time.
   b. Once all the APs preload for each managed device, the next few APs are preloaded and this step is repeated till all the APs are preloaded.
5. Once all the APs are preloaded, a target managed device is selected and rebooted.
6. After the target managed device boots up, partition targeted to this managed device is selected and an AP move is initiated. Once all the APs are rebooted, the next partition is selected and an AP move is initiated.
7. Once all the partitions targeted to the targeted managed device are upgraded successfully, another target managed device is selected and step 6 is repeated till all the managed devices are rebooted successfully.

Following section describes how to configure Live upgrade and the limitations of live upgrade:

Configuring Live Upgrade

The following procedure describes how to upgrade the cluster:

1. Log in to a Mobility Master.
2. In the **Managed Network** node hierarchy, navigate to **Maintenance > Software Management**..
3. Select one or multiple clusters from the table. The table only lists the clusters and not the cluster members.



The table displays the Name of the cluster, the number of managed devices, the number of APs in the cluster, and the version of software running on the managed devices of the cluster.

4. In the **INSTALLATION SETTINGS > When** menu, select **Now**.
5. In the **Image File** section, specify the image file location, name, and the protocol to use. Enter values for the following parameters:
   - **Server IP address**—IP address of the server
   - **Image path**—Path of image file
   - **Software to Install**—Version to upgrade to
   - **Protocol**—Protocol to use to transfer file. Valid values are: FTP, TFTP, and SCP. Default value is FTP.
   - **Username**—Username of the account on the image server.
   - **Password**—Password of the account on the server.

6. Click **Next.**
7. In the **Target Partition** section, specify the managed device partition where you want to install the firmware and where you want it to boot from. Select any one of the following partitions:
   - **Partition 0**
   - **Partition 1**
   - **auto**
8. Click **Install** to start the cluster upgrade process.
9. To check the status of the upgrade, in the **Managed Network** node hierarchy, navigate to **Configuration > Services > Cluster**.
10. The **upgrade status for cluster name** section is displayed with the status of the upgrade. The status of the upgrade can be any of the following values:
    - **Upgrade pending**
    - **AP partitioning in progress**
    - **Upgrade in progress**
    - **Upgrade failed**
    - **Upgrade completed**
11. The status of each managed device is displayed. Similarly, the status for each AP in the cluster is also displayed.
12. You can also upgrade a cluster using the shortcut icon by navigating to **Configuration > Services > Cluster** and selecting the cluster you want to upgrade.

Execute the following command to trigger the cluster upgrade in the Mobility Master:

```
(host) [mm] [cluster1] #lc-cluster <cluster_name> initiate upgrade version <img_
version> partition <partition_id>
```

- cluster_name: The configured cluster profile name; corresponds to the managed devices and APs associated to the cluster that needs to be upgraded.
- img_version: The target image version, such as 8.2.1_XXXXX.
- partition_id: The partition on the managed device to which the new image is to be copied, valid values are 0 or 1 and this is optional. If the partition is not specified, it automatically picks the alternate boot partition.



The fileserver is configured in upgrade profile. Ensure that the **upgrade profile** is configured before upgrading the cluster using the CLI.

To check the status of the upgrade, for Cluster1, execute the following show commands:

```
(host)#show lc-cluster cluster1 controller details
(host)#show lc-cluster cluster1 ap details
(host)#show lc-cluster cluster1 upgrade status
(host)#show lc-cluster cluster1 upgrade status verbose
(host)#show lc-cluster cluster1 upgrade stats
```

Limitation

This feature has the following limitations:

- As there is an image preload limitation, cluster upgrade cannot be done with two different versions without reloading the managed devices. Only one preload per managed device reboot is allowed currently.
- The state of the AP preload is not displayed in the WebUI during the cluster upgrade.
- Cluster upgrade cannot be triggered if a previous upgrade resulted in split subclusters.
- Cluster Rolling Upgrade is not supported on mesh nodes. Therefore, a cluster upgrade cannot be performed for mesh deployments.
- It is not recommended to use live upgrade in MultiZone deployments. Each zone can use live upgrade but the upgrade on datazone is not hitless.

Scheduling a Cluster Upgrade

Starting from ArubaOS 8.4.0.0, scheduled cluster upgrade allows you to schedule the upgrade to a specified time to avoid manual intervention. The cluster is upgraded automatically at the scheduled time. You can view, delete, or reschedule the scheduled cluster upgrade.

You can also schedule cluster upgrade for one or more profiles at the same time or different times.

When a Mobility Master reboots or a process restarts, this feature has the ability to preserve the configured scheduled upgrades. It also allows you to synchronize the scheduled upgrades to a standby Mobility Master. If a Mobility Master has active and standby devices configured, then the scheduled upgrade information will be synchronized between active and standby through database synchronization. The upgrade is initiated when the Mobility Master becomes active.

Key Considerations

- Upgrades can be schedule only to a future time, maximum of 30 days from the managed device's current time.
- The time scheduled is always in reference to the managed devices in a cluster.
- All the network nodes have to be NTP synchronized.

Limitation

- DST time change hour is not automatically adjusted for a scheduled upgrade.
- Time changed manually in a managed device is not automatically adjusted for a scheduled upgrade.

Configuring a Scheduled Upgrade

To configure a scheduled upgrade, perform the following steps either through WebUI or CLI:

To configure a scheduled cluster upgrade, refer to the Scheduling Upgrade of Clusters on page 966 section.

To schedule a scheduled cluster upgrade:

```
(host) [mm] (config) #lc-cluster <cluster_prof> schedule upgrade <version> <year> <month> <day> <hh> <mm> <ss>
```

Parameter	Description
`cluster_prof`	Cluster profile for which upgrade is scheduled
`version`	The version to which the cluster will get upgraded to
`year`	Year of the upgrade

Parameter	Description
month	Month of the upgrade
day	Day of the upgrade
hh	Hour of the upgrade
mm	Minutes of the upgrade
ss	Seconds of the upgrade

Example:

```
(host) [mm] (config) #lc-cluster <cluster_prof> schedule upgrade version 8.4.0.0-
sangiovese_73823 2018 04 10 00 00 00
```

Viewing the Scheduled Cluster Upgrade Status

```
(host) [mm] #show lc-cluster scheduled-upgrades
```

Example:

```
(host) [mm] #show lc-cluster scheduled-upgrades
Cluster Scheduled Upgrade Status
--------------------------------
Profile   To Version            Partition ID   AP Preload size   Scheduled Time
     MD Timezone
-------   ----------            ------------   ---------------  --------------
     -----------
v4         8.4.0.0-mm-dev_65200 Default        100               Fri Jun  8 15:00:00
2018   Asia/Tokyo
```

Deleting or Aborting a Scheduled Cluster Upgrade

```
(host) [mm] (config) #lc-cluster <cluster_prof> abort scheduled-upgrade
```

Rescheduling a Cluster Upgrade

```
(host) [mm] (config) #lc-cluster v4 re-schedule upgrade <version> <year> <month> <day>
<hh> <mm> <ss>
```



To reschedule a cluster upgrade, the upgrade must already be scheduled.

Example:

```
(host) [mm] (config)#lc-cluster v4 re-schedule upgrade version 8.2.0.1 2018 6 6 0 50 0
```

Troubleshooting Cluster

This section provides commands that can be used to troubleshoot different scenarios in a cluster configuration.

The different control plane processes in the cluster are GSM manager (GSM), cluster manager (CM), Station Manager (STM), and AUTH. On the AP, the main modules are A-STM and ASAP (datapath).

The following is a list of some common troubleshooting scenarios in a cluster:

- Cluster Formation Unsuccessful
- AP Rebootstrap
- Users are Unable to Connect to a Cluster
- Users are Getting Deauthenticated

Cluster Formation Unsuccessful

All managed devices in a cluster are collectively known as cluster members. The cluster formation is successful when all the managed devices in the cluster are connected to each other.

Some of the reasons because of which a cluster formation is unsuccessful are as follows:

1. If the cluster group membership is not executed.
2. If all the managed devices are not listed in cluster.
3. If there is a connectivity issue and managed devices are not able to reach their peer.
4. If IPsec SA is not formed.

To check the status of the cluster formation, execute the **show lc-cluster group membership** command.

```
(host) [mynode] #show lc-cluster group-membership
Mon Dec 21 17:30:51.952 2015
Cluster Enabled, Profile Name = "6NodeCluster"
Redundancy Mode On
Active Client Rebalance Threshold = 50%
Standby Client Rebalance Threshold = 75%
Unbalance Threshold = 5%
Cluster Info Table
------------------
Type IPv4 Address    Priority Connection-Type STATUS
---- --------------- -------- --------------- ------
self     10.15.116.3      128             N/A ISOLATED (Leader)
peer     10.15.116.4      128    L3-Connected CONNECTED-FROM-SELF-DISCONNECTED-FROM-
PEERS
peer     10.15.116.5      128    L3-Connected CONNECTED-FROM-SELF-DISCONNECTED-FROM-
PEERS
peer     10.15.116.8      128    L3-Connected CONNECTED-FROM-SELF-DISCONNECTED-FROM-
PEERS
peer     10.15.116.9      128             N/A SECURE-TUNNEL-NEGOTIATING
peer    10.15.116.10      128             N/A SECURE-TUNNEL-NEGOTIATING

DISCONNECTED
INCOMPATIBLE
DISCONNECTED-FROM-SELF-CONNECTED-FROM-PEERS",
CONNECTED-FROM-SELF-DISCONNECTED-FROM-PEERS",
SECURE-TUNNEL-NEGOTIATING
SECURE-TUNNEL-ESTABLISHED
CONNECTED
```

Table 71: *Cluster state*

State	Reason
INCOMPATIBLE	This error can occur in the following scenario: If two managed devices are running different ArubaOS versions, then a build string mismatch is found and the managed devices are not part of the cluster.
DISCONNECTED	This error can occur in the following scenario:

State	Reason
	■ If none of the managed devices in the cluster are in the CONNECTED state. ■ If there is an issue with the physical connectivity among the managed devices in the cluster. ■ If one of the ports is an untrusted node.
SECURE TUNNEL NEGOTIATION	This status is displayed for a very short period of time till the IPsec tunnel is set up. If the status persists, it indicates that there is an issue in the IPsec tunnel setup.
CONNECTED FROM SELF DISCONNECTED FROM PEER	This error can occur in the following scenario: ■ Managed device 1 and managed device 2 are connected. Managed Device 3 is later introduced in the cluster. Managed device 1 and managed device 3 are connected but managed device 2 and managed device 3 are not connected.

After the cluster moves to the CONNECTED state, check if it is L2-connected, where every VLAN on the peer is reachable as determined by VLAN probing. Use the following command to check the VLAN probing status:

```
(host) [mynode] #show lc-cluster vlan-probe status
```

Execute the VLAN probing algorithm on the managed device, if you have made some VLAN changes to the distribution switch:

```
(host) [mynode] (config) #lc-cluster start-vlan-probe
```

AP Rebootstrap

An AP rebootstraps when a S-AAC is not assigned to it. The following is a list of some reasons because of which an AP rebootstraps:

1. Platform capacity — If the managed device has reached its maximum capacity or it already has the maximum APs it can support.

   To resolve this issue, perform the following steps:

   - Add another managed device or upgrade an existing managed device to support more number of APs.
   - Rework on the network configuration.

2. Multiple managed devices are down – If an S-AAC goes down, the Standby Controller (S-UAC) is made the Active Controller (A-UAC). However, if the S-UAC also goes down, then the AP rebootstraps.

   To resolve this issue, ensure that you make an appropriate selection of the distribution switch to handle the required scale.

Users are Unable to Connect to a Cluster

The following is a list of some reasons why a user might be unable to connect to a cluster:

1. The AP and the managed device have different roles for the user.

   Every user has an A-UAC and if the AP's information of the UAC for a user is different from the actual managed device's information and if the managed device does not have this information regarding the user, then it rejects the user.

2. IPsec tunnel is not established.

   If CPsec is enabled on the APs, then the APs are expected to have the IPsec tunnel established with all the managed devices in the cluster. If the IPsec tunnel is not established, the user cannot connect to the cluster.

3. There is incomplete AP configuration for an 802.1X client.

   For 802.1X clients to connect, multicast key (mkey) has to go from the AAC to the UAC. If the mkey is not available in the UAC, the status is not displayed and the user cannot connect. To check for incomplete AP configurations, execute the **show auth-tracebuf** command .

Users are Getting Deauthenticated

The following is a list of some reasons why a user might get deauthenticated:

1. Cluster failover — If a user is deauthenticated in a cluster, check if there is a cluster failover at the same time. To check when a managed device in DOWN status was first disconnected, use the **show lc-cluster heartbeat counters** command.
   a. In case a failover occurs when the managed devices are down, check if the managed devices are L2-connected using the **show lc-cluster vlan-probe status** command.
   b. If the managed devices are L3-connected, fix the VLAN probe using the **lc-cluster exclude-vlan <vlan-number>** command.
2. If the managed devices are L2-connected and if the issue persists, check for a solution in AP Rebootstrap.
3. If the AP does not rebootstrap and if there is no fail over, contact Technical Support Team.

Enabling Debug

In a cluster setup, a lightweight tracing mechanism is added to collect debug information with minimal performance impact on the cluster.

In a 7200 Series managed device, the debug information gets collected in the flash1 partition of the managed device and can be used for future troubleshooting. In a 7000 Series and 7205 managed devices, there is no flash1 partition and a USB device is needed to collect this debug information, which can be used for future debugging or reporting of an issue.

Execute the following trace commands to collect debug information for the cluster:

```
(host) #gsm trace channel ap application stm
(host) #gsm trace channel ap application dds
(host) #gsm trace channel ap application cluster_mgr
(host) #gsm trace channel radio application stm
(host) #gsm trace channel radio application dds
(host) #gsm trace channel sta application stm
(host) #gsm trace channel sta application auth
(host) #gsm trace channel sta application dds
(host) #gsm trace channel sta application cluster_mgr
(host) #gsm trace channel mac_user application auth
(host) #gsm trace channel mac_user application dds
(host) #gsm trace channel mac_user application cluster_mgr
(host) #gsm trace channel ip_user application auth
(host) #gsm trace channel ip_user application dds
(host) #gsm trace channel user application auth
(host) #gsm trace channel user application dds
(host) #gsm trace channel sectun application dds
(host) #gsm trace channel sectun application cluster_mgr
(host) #gsm trace channel key_cache application auth
```

```
(host) #gsm trace channel key_cache application dds
(host) #gsm trace channel pmk_cache application stm
(host) #gsm trace channel pmk_cache application auth
(host) #gsm trace channel pmk_cache application dds
(host) #gsm trace channel rep_key application dds
(host) #gsm trace channel rep_key application cluster_mgr
(host) #gsm trace channel cluster application dds
(host) #gsm trace channel cluster application cluster_mgr
(host) #gsm trace channel bucket_map application stm
(host) #gsm trace channel bucket_map application auth
(host) #gsm trace channel bucket_map application dds
(host) #gsm trace channel bucket_map application cluster_mgr
(host) #gsm trace channel cluster_bss application dds
(host) #gsm trace channel cluster_bss application cluster_mgr
(host) #gsm trace channel cluster_aac application dds
(host) #gsm trace channel cluster_aac application cluster_mgr
(host) #gsm trace channel cluster_ap application dds
(host) #gsm trace channel cluster_ap application cluster_mgr
(host) #gsm trace channel bss application stm
(host) #gsm trace channel bss application auth
(host) #gsm trace channel bss application cluster_mgr
(host) #dds trace receive channel sta peer $peerIP
(host) #dds trace transmit channel sta peer $peerIP
(host) #dds trace receive channel ip_user peer $peerIP
(host) #dds trace transmit channel ip_user peer $peerIP
(host) #dds trace receive channel mac_user peer $peerIP
(host) #dds trace transmit channel mac_user peer $peerIP
(host) #dds trace receive channel key_cache peer $peerIP
(host) #dds trace transmit channel key_cache peer $peerIP
(host) #dds trace receive channel pmk_cache peer $peerIP
(host) #dds trace transmit channel pmk_cache peer $peerIP
(host) #dds trace receive channel bucket_map peer $peerIP
(host) #dds trace transmit channel bucket_map peer $peerIP
(host) #dds trace receive channel cluster_bss peer $peerIP
(host) #dds trace transmit channel cluster_bss peer $peerIP
(host) #dds trace receive channel cluster_sta peer $peerIP
(host) #dds trace transmit channel cluster_sta peer $peerIP
(host) #dds trace receive channel cac_usage peer $peerIP
(host) #dds trace transmit channel cac_usage peer $peerIP
(host) #dds trace receive channel cluster_aac peer $peerIP
(host) #dds trace transmit channel cluster_aac peer $peerIP
(host) #dds trace receive channel cluster_ap peer $peerIP
(host) #dds trace transmit channel cluster_ap peer $peerIP
(host) #ap debug stm-trace category all loglevel  debug
(host) #aaa auth-trace loglevel debug
(host) #scm intiate audit <peerip>
```

The **Dashboard** page provides an enhanced visibility into your network to view and monitor various information of the devices in the network.

You can view the context sensitive help for each field in the **Dashboard** UI by clicking the **help** link at the top-right corner of the WebUI. The field for which the help is defined appears as green. You can turn off the help by clicking **Done**.

The ArubaOS 8.0.0.0 adds the Mobility Master and managed device in a topology that permits the Mobility Master to manage and monitor one or more managed devices. For this, the administrator has to configure mgmt-server as the Mobility Master from the managed device's path. The Mobility Master can only manage and monitor managed devices and cannot manage APs.

In the **Managed Network** node hierarchy, navigate to **Dashboard**. You can view the **Dashboard** of a managed device without logging out of the Mobility Master. To view the **Dashboard** page of a managed device, click the managed device. See Table 72 for dashboard pages available in managed network and managed device.

Dashboard Pages

Starting from ArubaOS 8.4.0.0, the dashboard page contains the following sub-categories:

- Overview
- Infrastructure
- Traffic Analysis
- Security
- Services

The following table shows the available dashboard pages :

Table 72: *Dashboard Pages in Various Views*

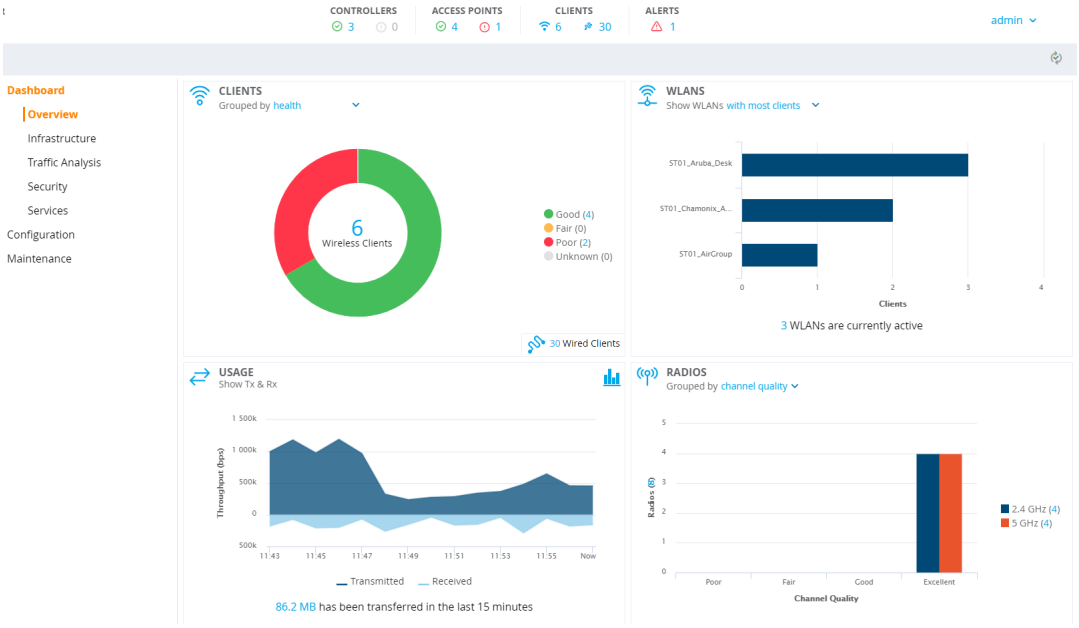
Dashboard Pages	Managed Network	Managed Device
Overview		
Clients	Yes	Yes
WLANs	Yes	Yes
Usage	Yes	Yes
Radios	Yes	Yes
Infrastructure		
Controller	Yes	Yes

Dashboard Pages	Managed Network	Managed Device
Access Devices	Yes	Yes
WAN	Yes	Yes
Cluster	Yes	—
Traffic Analysis		
Applications	Yes	Yes
Destinations	Yes	Yes
Hosts	Yes	Yes
Websites	Yes	Yes
Security		
Detected Radios	Yes	—
Detected Clients	Yes	—
Events	Yes	—
Blacklisted Clients	Yes	Yes
Services		
AirGroup Servers	Yes	Yes
AirGroup Clients	Yes	Yes
UCC	Yes	Yes

Overview

The **Overview** dashboard provides the summary of CLIENTS, WLANs, USAGE, and RADIOS. See Figure 36 for **Overview** page.

Figure 36 *Overview Page*



The **Overview** dashboard contains the following windows:

- **CLIENTS**—This window displays the information about all the clients connected to a managed device within the network hierarchy. The number of wired and wireless clients is displayed in the bottom right corner of the **CLIENTS** window. This is the default page. For more information, see Clients on page 414.

  You can view the following information using the **Grouped by** drop-down list.

  - **Health** — Displays the health score of the wireless clients connected to APs. The health score are Good, Fair, Poor, or Unknown. Click the donut chart area or hyperlinked number to navigate to the **Wireless clients** table and view the details of the wireless clients connected to APs.
  - **Band** — Displays the list of wireless clients under the 2.4 GHz and 5 GHz radio bands. Click the donut chart area or hyperlinked number to navigate to the **Wireless clients** table and view the details of the wireless clients connected to APs.
  - **Data Speed** — Displays the data speed (bps) of the wireless clients connected to APs. Click the vertical bar to navigate to the **Wireless clients** table and view the details of the selected wireless clients.
  - **Signal Quality** — Displays the SNR (dB) ranges of the wireless clients connected to APs. Click the vertical bar to navigate to the **Wireless clients** table and view the details of the selected wireless clients.
  - **Operating System** — Displays the number of clients that are running each type of operating systems. Click the OS area in the donut chart or hyperlinked number to navigate to the **Wireless clients** table and view the details of the wireless clients running the selected operating system type.

- **WLANS**—This window displays the summary of all active WLANs in the managed device. This is the default page. Click the horizontal bar or hyperlinked number to navigate to the **WLANs** table and view the details of the active WLANs. For more information, see WLANs on page 417.

  You can view the following information using the **Show WLANs** drop-down list.

  - **With most clients** — Displays five WLANs currently accessed by highest number of clients, in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the **WLANs** table

and view the details of the selected WLAN.

- **With highest usage** — Displays five WLANs with highest number of bytes transmitted and received, in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the **WLANs** table and view the details of the selected WLAN.

- **USAGE**—Displays the throughput data (bps) transmitted and received in the last 15 minutes. Click the hyperlinked number to navigate to the **Usage** page and view the details of the data transmitted and received. For more information, see Usage on page 418.

Click **Current** icon in the top right corner of the window to display the data transmitted and received in the last 15 minutes.

- **RADIOS**— Displays the information of all radios of APs controlled by the managed device. For more information, see Radios on page 420.

  You can view the following information using the **Grouped by** drop-down list.

  - **Channel Quality** — Displays the channel quality of the 2.4 GHz and 5 GHz radio bands. This is a default page. Click the vertical bar to navigate to the **Radios** table and view the details of the 2.4 GHz or 5 GHz radio bands.
  - **Interference** — Displays the percentage of interference in the 2.4 GHz and 5 GHz radio bands. Click the vertical bar to navigate to the **Radios** table and view the details of the 2.4 GHz or 5 GHz radio bands.
  - **Channel Busy** — Displays the percentage of busy channel in the 2.4 GHz and 5 GHz radio bands. Click the vertical bar to navigate to the **Radios** table and view the details of the 2.4 GHz or 5 GHz radio bands.
  - **Channel** — Displays the active channels on the 2.4 GHz or 5 GHz radio bands. Click the vertical bar to navigate to the **Radios** table and view the details of the individual radio bands.

Mesh Links Table

Navigate to **Dashboard > Overview** and click on **Mesh Links** in the **Radios** window.



The **Radios** window displays **Mesh Links** only when mesh links are available.

The **Mesh Links** table is displayed with the following information:

- Cluster name
- Status
- Uplink Age
- Portal Access Point
- Mesh Points
- Band

Expand a mesh link to view the information related to **Link Statuses, Usage and Signal** details.

The **Link Statuses** window displays the following information:

- **Portal:** Displays the status of AP Portal. Clicking on the AP name will re-direct the user to Access Points page.
- **Points:** Displays the status of all mesh points. Clicking on the AP name re-direct the user to Access Points page.
- **Uptime:** Displays uptime of the mesh portal.

- **SNR/Uplink age:** Hover over SNR/ Uplink Age to view the SNR values (in db) and the Uplink age. The color of the battery icon indicates the following:
  - Green- SNR value is 35db or higher
  - Orange- SNR value is between 20db to 35db
  - Red- SNR value is below 20db

The **Usage** window displays the **throughput** and **goodput** data for **all links** or for a specific mesh link.

- **Throughput** - The **Show throughput** displays the transmitted and received data (bps) in the last 15 minutes.
- **Goodput** - The **Show goodput** displays the goodput data (bps) in the last 15 minutes.

The **Signal** window displays the **channel utilization, transferred frames,** and **snr** for all links or for a specific mesh link.

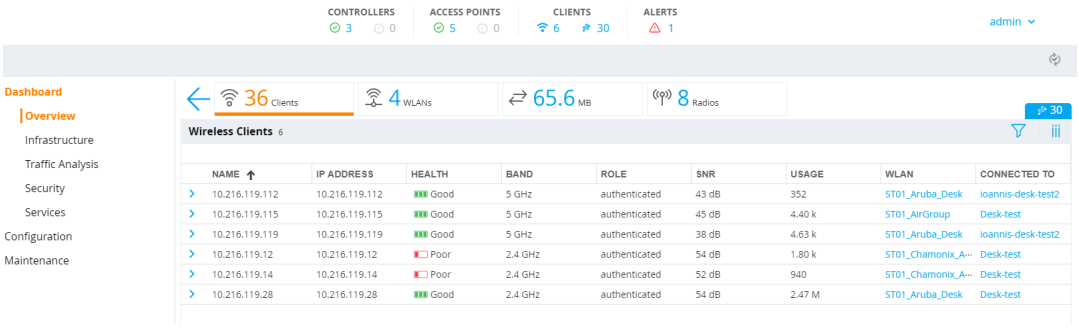
- **Channel utilization**- The **Show channel utilization** displays the percentage of the current channel utilization. The channel utilization information is categorized into Tx time, Rx time, Interference, and Free.
- **Transferred frames**- The **Show transferred frames** displays information about transmitted or received frames in the 2.4 GHz or 5GHz band. This information is categorized into Successful, Retried, and Dropped.
- **SNR**: The **Show snr** displays the SNR (db) ranges of mesh links.

Clients

The **Dashboard > Overview** page displays the details of both wired and wireless clients connected to the managed device. Click on **Wired Clients** or **Wireless Clients** to see the detailed list.

Starting from ArubaOS 8.7.0.0, the **Dashboard > Overview** page displays the details of wired clients connected in bridge mode. See Figure 37 for the **Clients** Page.

Figure 37 *Clients Page*



Clicking on the icon on top right-hand corner of the table will switch between the wired and wireless clients list. The information displayed in the **Role** column will provide details on the following:

- User-based tunneled users
- Port-based tunneled users
- Controller wired users
- AP wired users
- VIA-VPN users

Use the **Customize columns** option to choose the columns you want to view.

Action Bar

The Action bar displays the total number of wireless clients depending on the filters applied. The Action bar includes action buttons namely, **Show/Hide table filters**, **Add to blacklist**, and **Customize columns**.

NOTE

Aruba now allows you to manage blacklisted clients in stand-alone controllers as well as in a Mobility Master-Managed Device topology.

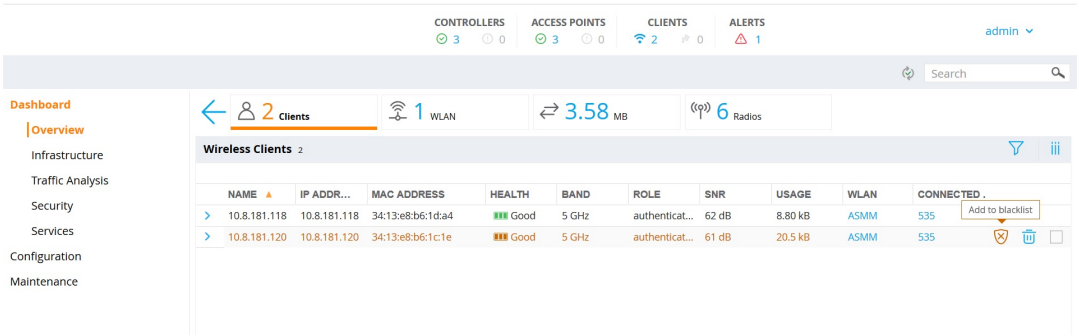
You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the wireless clients in the table that you want to view.
- **Add to blacklist** — Click the **Add to blacklist** icon to blacklist a wireless client.
- **Customize columns** — Click **Customize columns** icon to select the table columns that you want to view.

To blacklist a client manually, perform the following steps:

1. Select a client from the **Wireless Clients** table.
2. Click **Add to blacklist** icon as shown in Figure 38.

   Figure 38 *Add to Blacklist Page*

   The **Add to Blacklist** pop-up window is displayed.

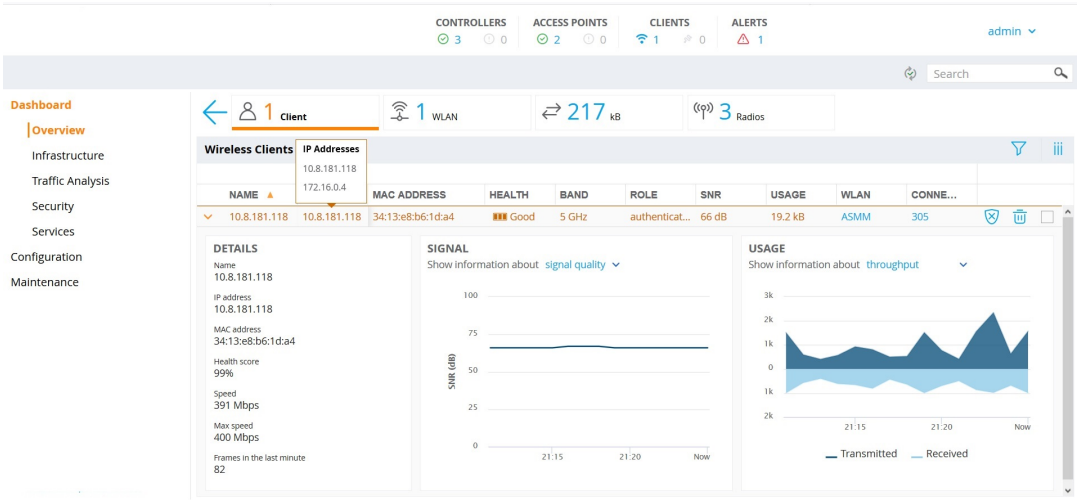
3. In the **Add to Blacklist** pop-up window, click **Add**.

   The client is blacklisted and is listed in the **Blacklisted Clients** table.

Details

Expand the wireless client from the **Wireless clients** table to view the detailed information of individual client. See Figure 39 for the Details Page.

Figure 39 *Details Page*



The **Wireless clients** table displays the following details:

- **Details** — Displays detailed information about the selected wireless client.
- **Signal** — Displays detailed information about throughput, transferred frames, signal quality, and data speed of the wireless clients.

You can view the following information using the **Show information about** drop-down list.

- **Throughput** — Displays the transmitted or received data (bps) in the last 15 minutes.
- **Transferred Frames** — Displays the information about transmitted or received frames in the 2.4 GHz or 5 GHz band. The transmitted or received frames information is categorized as: **Successful**, **Retried**, and **Dropped**.

Click **Historical** icon in the top right corner of the window to display transmitted or received frames in the last 15 minutes.

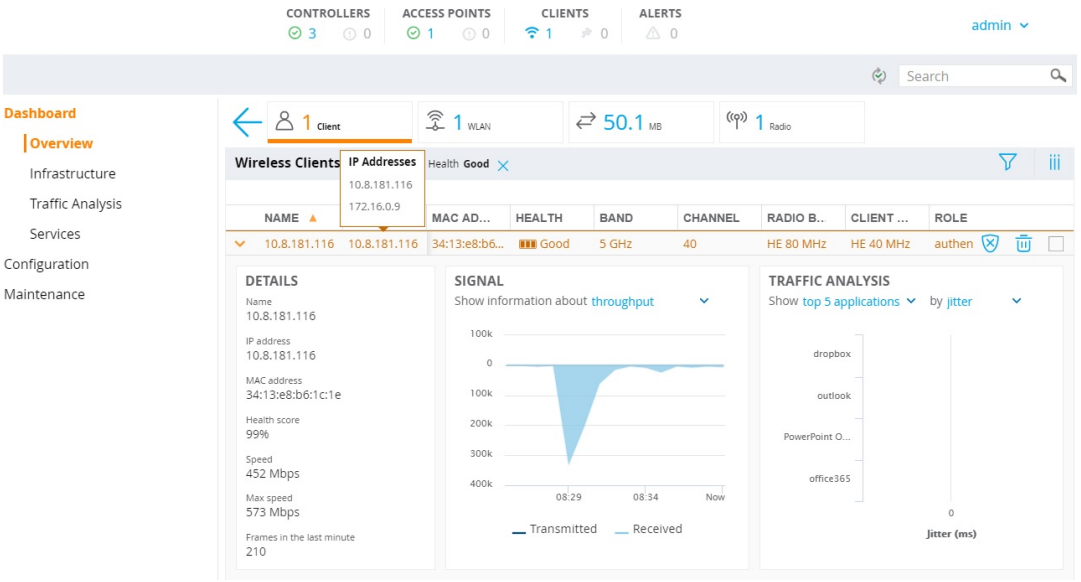
- **Signal Quality** — Displays the SNR (dB) ranges of the wireless clients in the last 15 minutes.
- **Data Speed** — Displays the data speed (bps) of the wireless clients in the last 15 minutes.
- **Traffic Analysis** — Displays detailed information about active applications and destinations used by the wireless clients.

You can view the following information using the **Show** drop-down list.

- **Top 5 Applications** — Displays five applications with highest usage (bytes), in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the **Applications** table and view the details of the wireless clients using this application.
- **Top 5 Destinations** — Displays five destinations with highest usage (bytes), in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the **Destinations** table and view the details of the wireless clients using this destination.

ArubaOS now allows you to view multiple IP addresses for a particular client. Hover your mouse over the IP address of a client to view a tooltip that lists all the IP addresses (IPv4 or IPv6 address, or a combination of both) associated with the client. See Figure 40 for the IP address list.

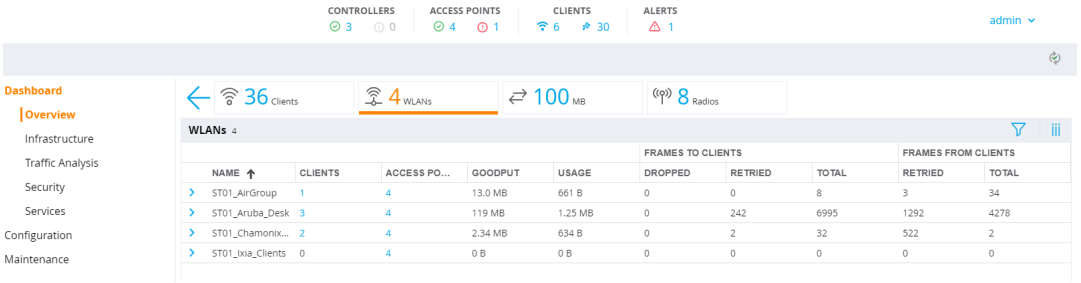
Figure 40 *Client IP Address List*



WLANs

Navigate to **Dashboard > Overview** and click **WLANS** icon. The **WLANs** page displays the summary of all the active WLANs in the managed device. See for **WLANs** page.

Figure 41 *WLANs page*



Action Bar

The Action bar displays the total number of wireless clients depending on the filters applied. The Action bar includes action buttons namely, **Show/Hide table filters**, and **Customize columns**.

You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the WLANs in the table that you want to view.
- **Customize columns** — Click **Customize columns** icon to select the table columns that you want to view.

Details

Expand the WLAN from the **WLANs** table to view the detailed information of individual WLAN. See Figure 42 for **Details page**.

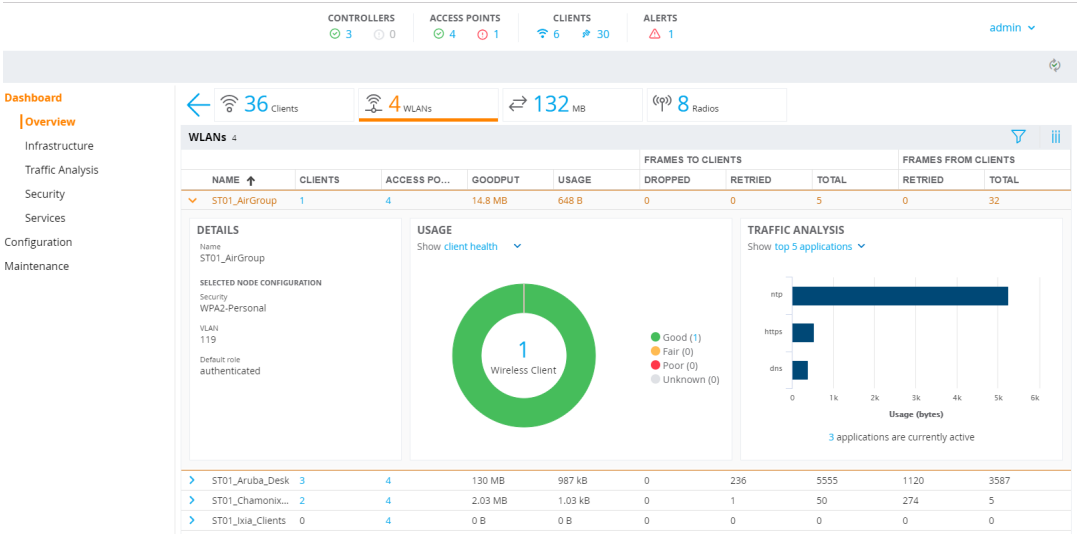
The **WLANs** table displays the following details:

- **Details** — Displays detailed information about the selected WLAN.
- **Usage** — Displays detailed information about health, band, and throughput of the wireless clients.

  You can view the following information using the **Show** drop-down list.
  - **Client Health** — Displays the health score of the wireless clients connected to APs. The health score are Good, Fair, Poor, or Unknown. Click the donut chart area or hyperlinked number to navigate to the **Wireless clients** table and view the details of the wireless clients connected to APs.
  - **Client by band** — Displays the list of wireless clients under the 2.4 GHz and 5 GHz radio band. Click the donut chart area or hyperlinked number to navigate to the **Wireless clients** table and view the details of the wireless clients connected to APs.
  - **Throughput** — Displays the transmitted or received data (bps) in the last 15 minutes.
- **Traffic Analysis** — Displays detailed information about active applications, destinations, clients and clients by OS.

  You can view the following information using the **Show** drop-down list.
  - **Top 5 Applications** — Displays five applications with highest usage (bytes), in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the **Applications** table and view the details of the wireless clients using the application.
  - **Top 5 Destinations** — Displays five destinations with highest usage (bytes), in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the **Destinations** table and view the details of the wireless clients using the destination.
  - **Top 5 Clients** — Displays five clients with highest usage (bytes), in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the **Hosts** table and view the details of the hosts.
  - **Clients by OS** — Displays the number of hosts that are running each type of OS. Click the OS area in the donut chart to display the hosts running the selected OS type.

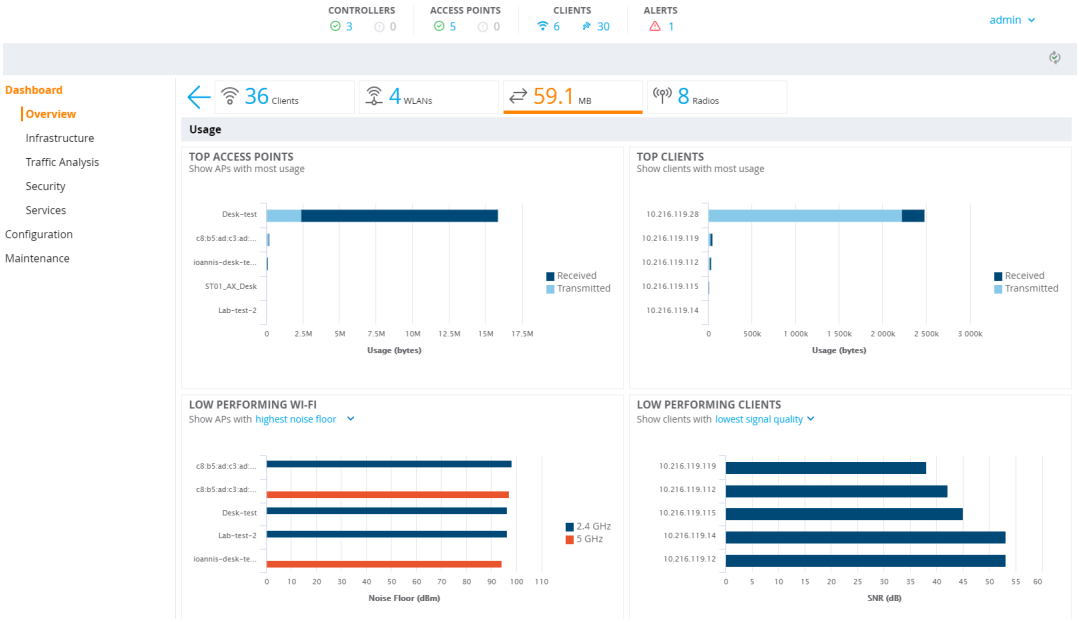
Figure 42 *Details Page*



Usage

Navigate to **Dashboard > Overview** and click **USAGE** icon. The **Usage** page displays the usage summary of APs, Clients, Low Performing Wi-Fi, and Low Performing Clients on the managed devices in the network. See Figure 43 for **Usage** page.

Figure 43 *Usage Page*



TOP ACCESS POINTS — Displays five access points with highest number of bytes transmitted or received, in decreasing order. Click the horizontal bar to navigate to the **Access points** table and view the details of the selected access point. For more information, see Details.

TOP CLIENTS — Displays five clients with highest number of bytes transmitted or received, in decreasing order. Click the horizontal bar to navigate to the **Wireless clients** table and view the details of the selected wireless client. For more information, see Details.

LOW PERFORMING WI-FI — Displays the details of low performing access points in the network. You can view the following information using the **Show APs with** drop-down list.

- **Highest noise floor** — Displays five access points with noise floor (dBm) in the 2.4 GHz or 5 GHz band. Click the horizontal bar to navigate to the **Access points** table and view the details of the selected access point. For more information, see Details.
- **Busiest Channel** — Displays five access points with percentage of busy channel in the 2.4 GHz or 5 GHz band. Click the horizontal bar to navigate to the **Access points** table and view the details of the selected access point. For more information, see Details.
- **Highest Interference** — Displays five access points with percentage of interference in the 2.4 GHz or 5 GHz band. Click the horizontal bar to navigate to the **Access points** table and view the details of the selected access point. For more information, see Details.

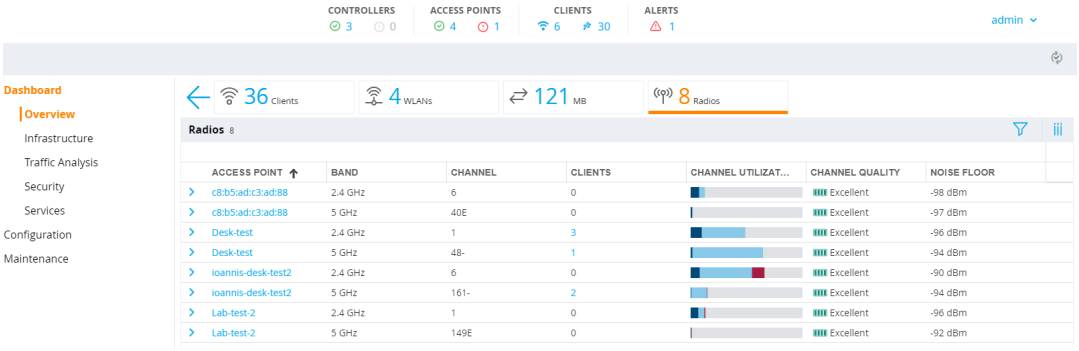
LOW PERFORMING CLIENTS — Displays the details of low performing access points in the network. You can view the following information using the **Show clients with** drop-down list.

- **Lowest signal quality** — Displays five wireless clients with lowest SNR (dB) ranges. Click the horizontal bar to navigate to the **Wireless clients** table and view the details of the selected wireless client. For more information, see Details.
- **Lowest Goodput** — Displays five wireless clients with lowest goodput (bps). Click the horizontal bar to navigate to the **Wireless clients** table and view the details of the selected wireless client. For more information, see Details.
- **Lowest data speed** — Displays five wireless clients with lowest data speed (bps). Click the horizontal bar to navigate to the **Wireless clients** table and view the details of the selected wireless client. For more information, see Details.

Radios

Navigate to **Dashboard > Overview** and click **RADIOS** icon. The **Radios** page displays the summary of all the active radios in the managed device. See Figure 44 for the **Radio** page.

Figure 44 *Radios Page*



Action Bar

The Action bar displays the total number of radios depending on the filters applied. The Action bar includes action buttons namely, **Show/Hide table filters**, and **Customize columns**.

You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the radios in the table that you want to view.
- **Customize columns** — Click **Customize columns** icon to select the table columns that you want to view.

Details

Expand the radios from the **Radios** table to view the detailed information of individual radios. See Figure 45 for Details page.

The **Radios** table displays the following details:

- **DETAILS** — Displays detailed information about the selected radio.
- **USAGE** — Displays detailed information about client health, connected clients, throughput, and goodput of the radio.

You can view the following information using the **Show information about** drop-down list.

- **Client Health** — Displays the health score of the wireless clients connected to APs. The health score are Good, Fair, Poor, or Unknown. Click the donut chart area or hyperlinked number to navigate to the **Wireless clients** table and view the details of the wireless clients connected to APs.
- **Clients Connected** — Displays the number of clients connected in the last 15 minutes.
- **Throughput** — Displays the transmitted and received data (bps) in the last 15 minutes.
- **Goodput** — Displays the goodput data (bps) in the last 15 minutes.
- **Channel** — Displays the detailed information about channel utilization, noise floor, transmitted or received frames, and WLANs.

You can view the following information using the **Show information about** drop-down list.

- **Channel Utilization** — Displays the percentage of the current channel utilization. The channel utilization information is categorized as: **Tx time**, **Rx time**, **Interference** and **Free**.

Click **Historical** icon in the top right corner of the window to display the channel utilization percentage in the last 15 minutes.

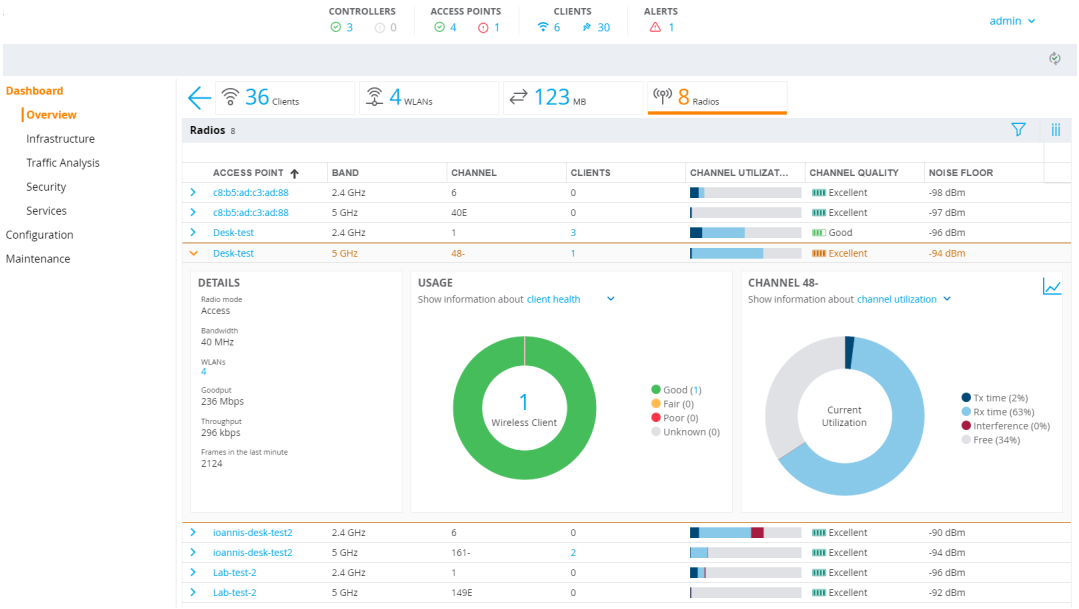
- **Noise Floor** — Displays the information about noise floor (dBm) in the last 15 minutes.
- **Transferred Frames** — Displays the percentage of transmitted or received frames in the channel. The transmitted or received frame information is categorized as: **Successful**, **Retried**, and **Dropped**.

Click **Historical** icon in the top right corner of the window to display the transmitted or received frames in the last 15 minutes.

Click **Clients** button in the bottom right corner of the window to navigate to the **Wireless clients** table and view the details of the wireless clients.

- **WLANs** — Displays the details of the throughput data (bps) of the selected channel. Click the horizontal bar to navigate to the **Wireless clients** table and view the details of the wireless clients connected to this WLAN. For more information, see Details.

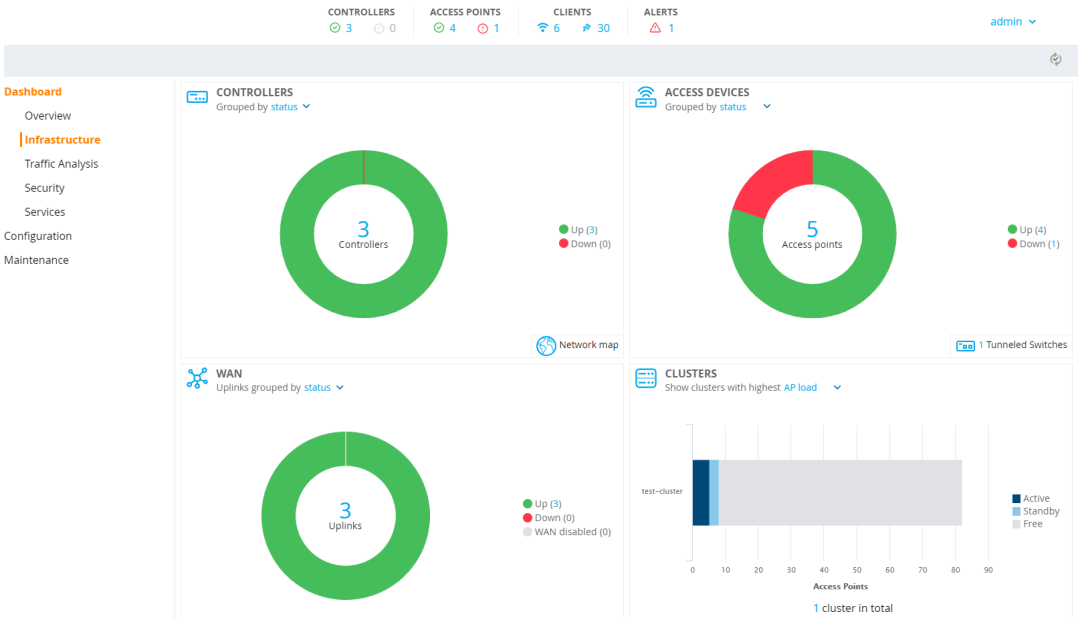
Figure 45 *Details Page*



Infrastructure

The **Infrastructure** provides the summary of Controllers, Access devices, WAN and Clusters. See Figure 46 for **Infrasructure** page.

Figure 46 *Infrastructure Page*



The **Infrastructure** dashboard contains the following windows:

- **CONTROLLERS**—This window displays the summary of all the controllers in the network. This is the default page. Click the donut chart or hyperlinked number to navigate to the **Controllers** table and view the details of all the controllers. For more information, see Controller on page 423.

You can view the following information using the **Grouped by** drop-down list.

- **Status** — Displays the status of all the controllers in the network. The controller status is categorized as: **Up**, and **Down**
- **Health** — Displays the health of all the controllers in the network. The controller health is categorized as:
- **Poor**- The health status is displayed as **Poor** when any one of the following scenarios is observed:
  - the controller itself is down.
  - at least one of its uplinks has poor health score.
  - 10% or more of the APs are Down.
- **Fair** - The health status is displayed as **Fair** when when any one of the following scenarios is observed:
    - at least one of its uplinks has a fair health score.
    - 1% or more (less than 10%) of the APs are Down.
  - **Good**- The health status is displayed as **Good** if APs and WAN uplinks are up.
  - **Unknown**- The health status is displayed as **Unknown** when the health status of the controller cannot be identified as good, fair, or poor.

Click **Network map** button to display the location of the managed devices in the network. If a managed device is not positioned, it is listed in a red balloon in the top-right corner of the map.

- **ACCESS DEVICES**—This window displays the summary of all the access points connected to a managed device. This is the default page. Click the donut chart or hyperlinked number to navigate to the **Access Points** table and view the details of all the access points connected to the network. For more information, see Access Devices on page 424.

You can view the following information using the **Grouped by** drop-down list.

- **Status** — Displays the status of all the access points connected to a managed device. The access point status is categorized as: **Up**, and **Down**.
- **AP Group** — Displays five access point groups with the number of access points connected per group.

Click **Tunneled Switches** button in the bottom right corner of the window to navigate to the **Tunneled Switches** table and view the details of tunneled switch.

- **WAN** — The window displays the summary of status and health of uplink in the network. This is the default page. Click the donut chart or hyperlinked number to navigate to the **Uplinks** table and view the details of all the uplinks in the network. For more information, see Details on page 427.

You can view the following information using the **Uplink grouped by** drop-down list.

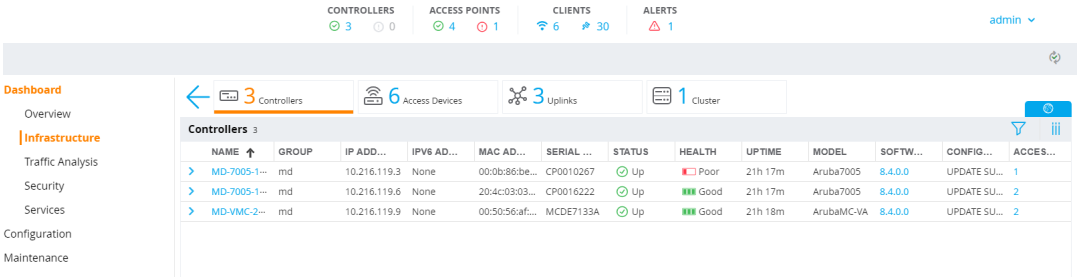
- **status** — Displays the status of all uplinks in the network. The uplink status is categorized as: **Up**, **Down** and **WAN disabled**
- **health** — Displays the health of all uplinks in the network. The uplink health is categorized as: **Good**, **Fair**, **Poor** and **Unknown**

- **CLUSTERS** — The **Cluster** dashboard provides a visual overview on each cluster deployed on the network. The cluster dashboard displays total AP load, client load per cluster, health of each cluster, and status of each controller and access point connected to the cluster.
- **Clients**—This window displays the summary of all active wired and wireless clients in the managed device. This is the default page. Click the horizontal bar or hyperlinked number to navigate to the **Clients** table and view the details of the active clients. For more information, see Details on page 415.

Controller

Navigate to **Dashboard > Infrastructure** and click **CONTROLLERS** icon. The **Controllers** page lists all the managed devices in the network and provides its status and health related information. See Figure 47 for **Controllers** page.

Figure 47 *Controller Page*



Action Bar

The Action bar displays the total number of controllers depending on the filters applied. The Action bar includes action buttons namely, **Show/Hide table filters**, and **Customize columns**.

You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the controllers in the table that you want to view.
- **Customize columns** — Click **Customize columns** icon to select the table columns that you want to view.

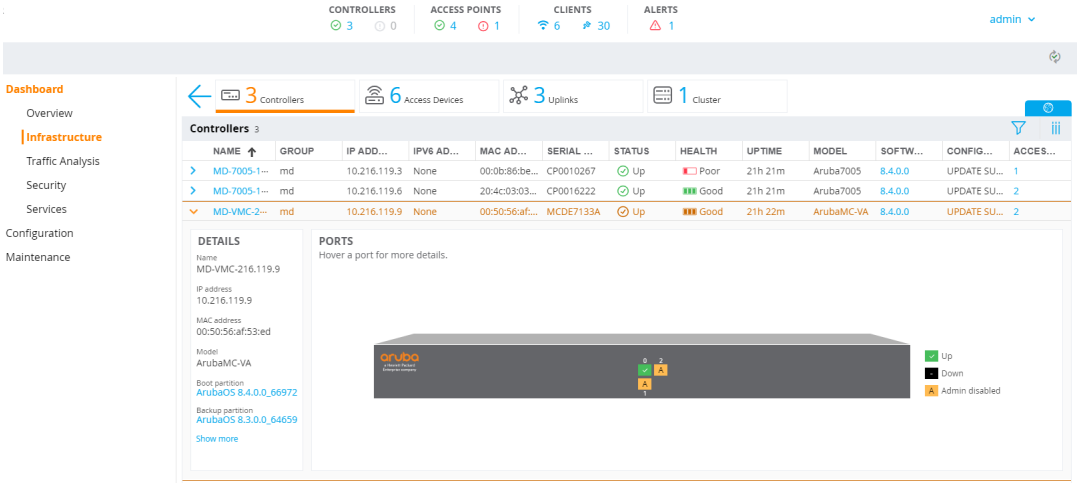
Details

Expand the controller from the **Controllers** table to view the detailed information of individual controller. See Figure 48 for Details page.

The **Controllers** table displays the following details:

- **Details** — Displays detailed information about the managed device.
- **Ports** — Displays the status of all the ports in the managed device.

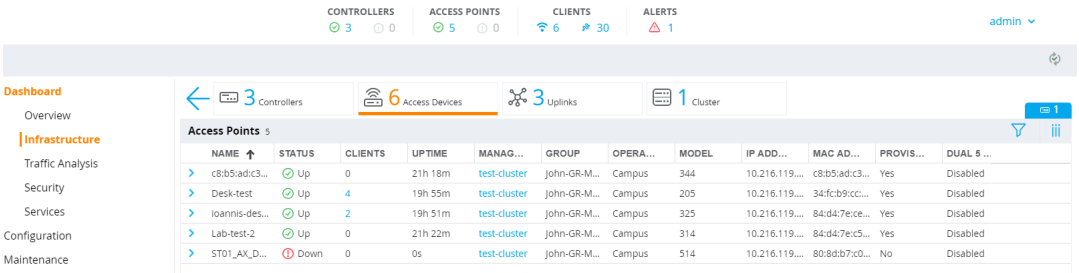
Figure 48 *Details Page*



Access Devices

Navigate to **Dashboard > Infrastructure** and click **ACCESS DEVICES** icon. The **Access Points** page lists all the access points connected to managed devices in the network and provides its status and access point group related information. See Figure 49 for **Access Points** page.

Figure 49 *Access Points Page*



Action Bar

The Action bar displays the total number of APs depending on the filters applied. The Action bar includes action buttons namely, **Show/Hide table filters**, and **Customize columns**.

You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the access points in the table that you want to view.
- **Customize columns** — Click **Customize columns** icon to select the table columns that you want to view.
- **Sort**: Click a column header of the Access Points table to sort the complete list based on the entries on the active column. You can also use the sort icon that appears when you click on a column for sorting.

- **View client details**: Click on the number of clients associated with the AP to view the details of the clients on the **Clients** page.
- **View AP or radio summary**: Expand the access point from the **Access Points** table to view the summary of the individual APs. See Details on page 425 for more information.
- **Capture and download packets:** Click Packet Capture icon of an access point from the Action column to start, pause, or stop capturing and downloading the AP packets.
- **Delete**: Select the check box for APs with **Down** status, and click the trash icon to remove the APs from the table. The APs with **Down** status are either unused or replaced when deployed.



You can delete the APs with **Down** status only on Mobility Master or standalone controllers, and not on managed devices.

Details

Expand a access point from the **Access Points** table to view the detailed information of individual access point. See Figure 50 for Details page.

The **Access Points** table displays the following details:

- **Details** — Displays detailed information about the selected access point.
- **Radio 2.4 GHz Channel** — Displays the about channel utilization, noise floor, transmitted or received frames, and WLANs in the 2.4 GHz band.

You can view the following information using the **Show information about** drop-down list.

- **Channel Utilization** — Displays the percentage of the current channel utilization in the 2.4 GHz band. The channel utilization information is categorized as: **Tx time**, **Rx time**, **Interference**, and **Free**.

Click **Historical** icon in the top right corner of the window to display the percentage of the current channel utilization in the last 15 minutes.

- **Noise Floor** — Displays the information about noise floor (dBm) in the last 15 minutes.
- **Transferred Frames** — Displays the information about transmitted or received frames in the 2.4 GHz band. The transmitted or received frame information is categorized as: **Successful**, **Retried**, and **Dropped**.

Click **Historical** icon in the top right corner of the window to display the transmitted or received frames in the last 15 minutes.

- **WLANs** — Displays the throughput data (bps) in the 2.4 GHz band. Click the horizontal bar to navigate to the **Wireless clients** table and view the details of the wireless clients connected to this WLAN. For more information, see Details.

Radio 5 GHz Channel — Displays the detailed information about channel utilization, noise floor, transmitted or received frames, and WLANs in the 5 GHz band.

You can view the following information using the **Show information about** drop-down list.

- **Channel Utilization** — Displays the percentage of the current channel utilization in the 5 GHz band. The channel utilization information is categorized as: **Tx time**, **Rx time**, **Interference**, and **Free**.

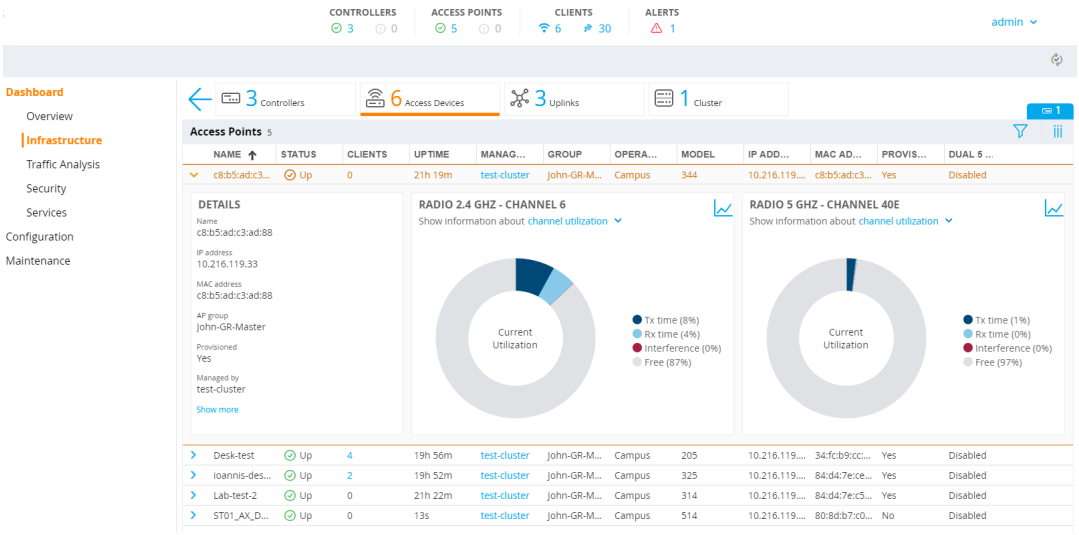
Click **Historical** icon in the top right corner of the window to display the channel utilization percentage in the last 15 minutes.

- **Noise Floor** — Displays the information about noise floor (dBm) in the last 15 minutes.
- **Transferred Frames** — Displays the information about transmitted or received frames in the 5 GHz band. The transmitted and received frame information is categorized as: **Successful**, **Retried**, and **Dropped**.

Click **Historical** icon in the top right corner of the window to display transmitted or received frames in the last 15 minutes.

- **WLANs** — Displays the throughput data (bps) in the 5 GHz band. Click the horizontal bar to navigate to the **Wireless clients** table and view the details of the wireless clients connected to this WLAN. For more information, see Details.

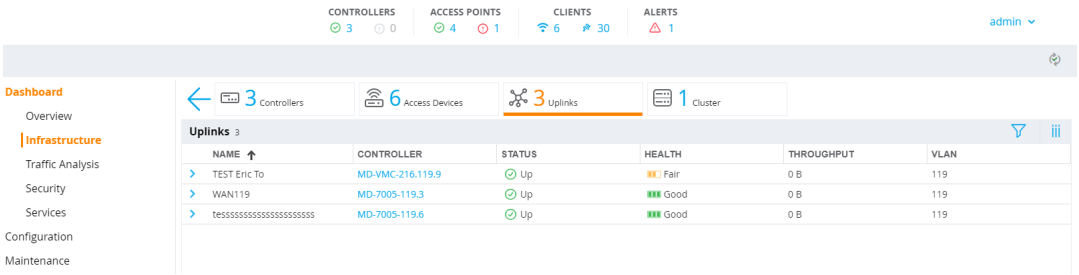
Figure 50 *Details Page*



WAN

Navigate to **Dashboard > Infrastructure** and click **WAN** icon. The **Uplinks** page provides the status and health related information of uplinks in the network. See Figure 51 for **Uplinks** page.

Figure 51 *Uplinks Page*



Action Bar

The Action bar displays the total number of uplinks depending on the filters applied. The Action bar includes action buttons namely, **Show/Hide table filters**, and **Customize columns**.

You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the uplinks in the table that you want to view.

- **Customize columns** — Click **Customize columns** icon to select the table columns that you want to view.

Details

Expand a uplink from the **Uplinks** table to view the detailed information of individual uplink. See Figure 52 for Details page.

The **Uplinks** table displays the following details:S

- **HEALTH** — Displays detailed information about jitter, latency, and health score of the uplink in the network.

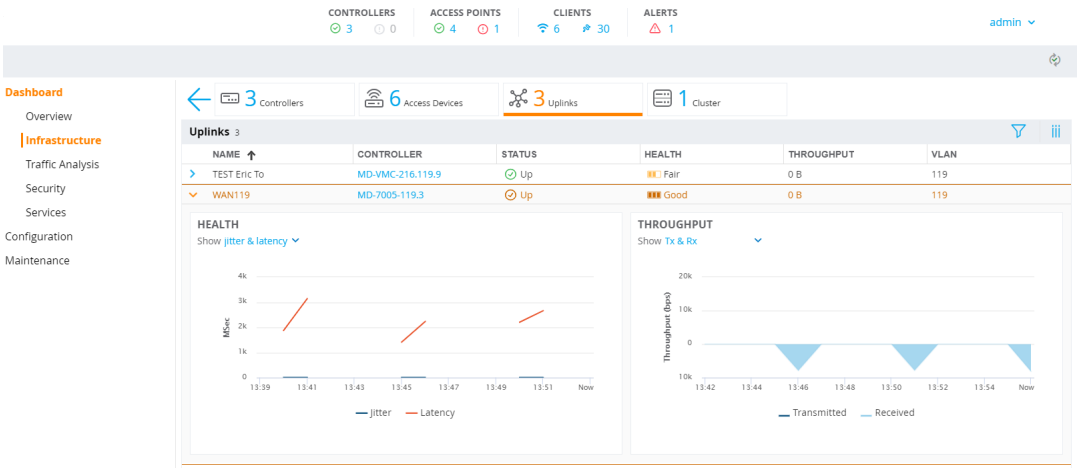
  You can view the following information using the **Show** drop-down list.

  - **Jitter and Latency** — Displays the jitter and latency (Msec) in the last 15 minutes.
  - **Health Score** — Displays the percentage of health score in the last 15 minutes.
- **THROUGHPUT** — Displays detailed information about transmitted or received data, and global compression on the uplink.

  You can view the following information using the **Show** drop-down list.

  - **Tx and Rx** — Displays the transmitted or received data (bps) in the last 15 minutes.
  - **Global Compression** — Displays the aggregated compression saving on every uplink of the controller in the last 15 minutes.

Figure 52 *Details Page*



Cluster

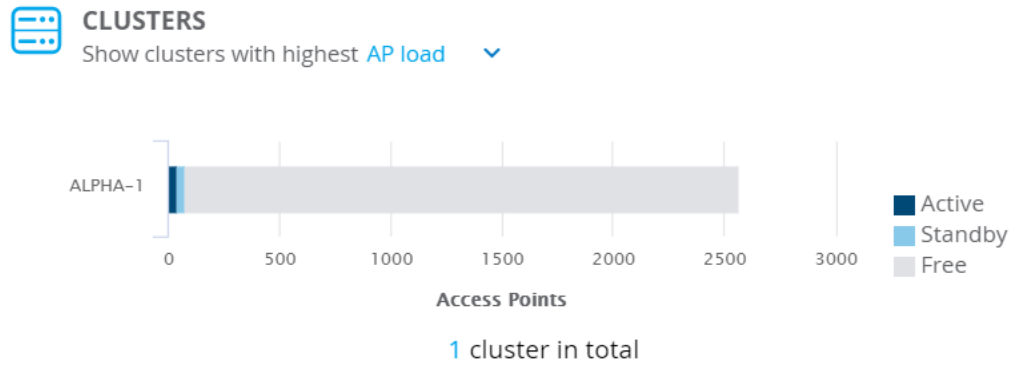
The **Cluster** dashboard provides a visual overview on each cluster deployed on the network, displaying the following information:

- Total AP load per cluster
- Total Client load per cluster
- Status of each controller and access point in a cluster
- Health of each cluster

Traditionally, cluster dashboard was visible only on the managed network. However, starting from ArubaOS 8.7.0.0, cluster dashboard can be accessed using all nodes. To view the **Cluster** dashboard, navigate to **Dashboard > Infrastructure > Clusters** in the WebUI. By default, the cluster dashboard displays the cluster with the highest AP load.

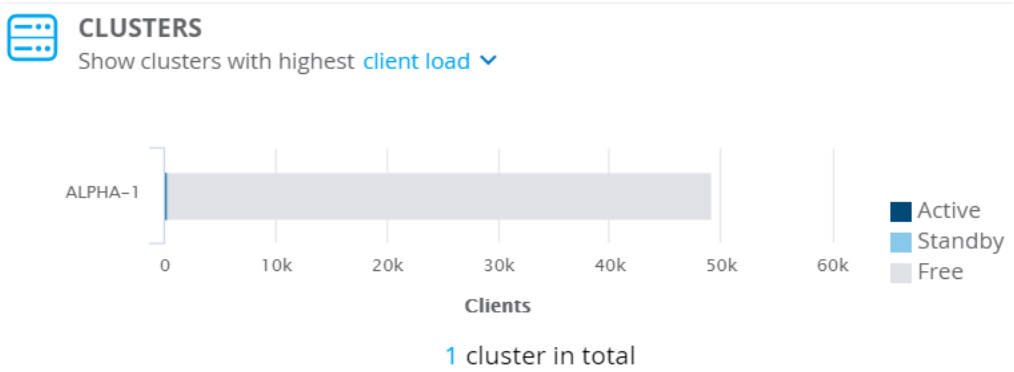
Following dashboard displays the cluster with the highest AP load. It also displays the number of active, standby and free APs for a given cluster.

Figure 53 *The Cluster Dashboard with Highest AP Load*



To view the client load, you can select client load to display the clusters with the highest client load from the drop-down menu.

Figure 54 *The Cluster Dashboard with Highest Client Load*



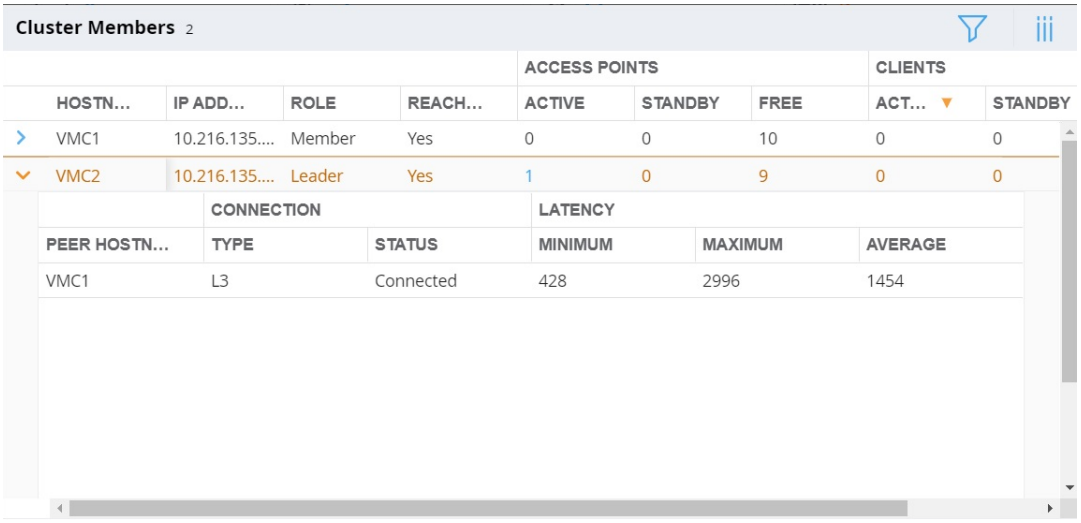
The **Cluster** dashboard consists of an **Cluster** section and **Cluster Member** section. Click on the AP load page or client load page to view more details.

- **Cluster > AP Load**: Displays the proportional distribution and number of active, standby, and free APs. Hover your mouse above a section of the chart to view the count for that AP type:
  - Free AP Load
  - Active AP Load
  - Standby AP Load
  - Total AP Load
- **Cluster > Client Load**: Displays the proportional distribution and number of active, standby, and free stations (clients). Hover your mouse above a section of the chart to view the count for that station type:
  - Free STA Load
  - Active STA Load
  - Standby STA Load
  - Total STA Load

To view in-depth information of each cluster member, click on the hyperlinked number under the **Controllers** column of the **Clusters** table. A **Cluster Members** pop-up window is displayed that

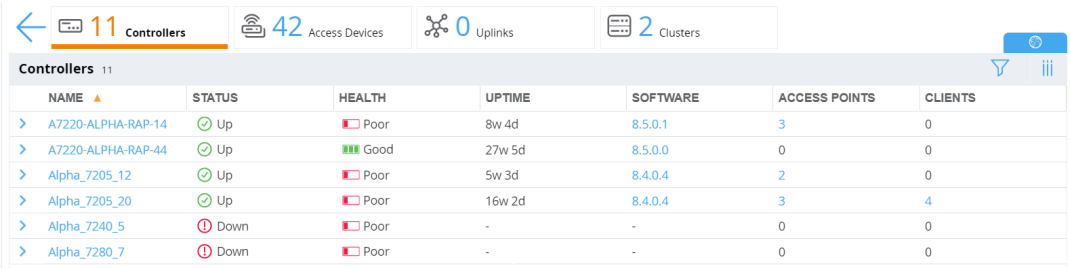
contains a summary of each cluster member such as hostname, IP address, the cluster roles and so on. You can expand a particular cluster member to view the connection status and type (L2, L3, or both) of that cluster member as shown in the following figure:

Figure 55 *Cluster Member Information*



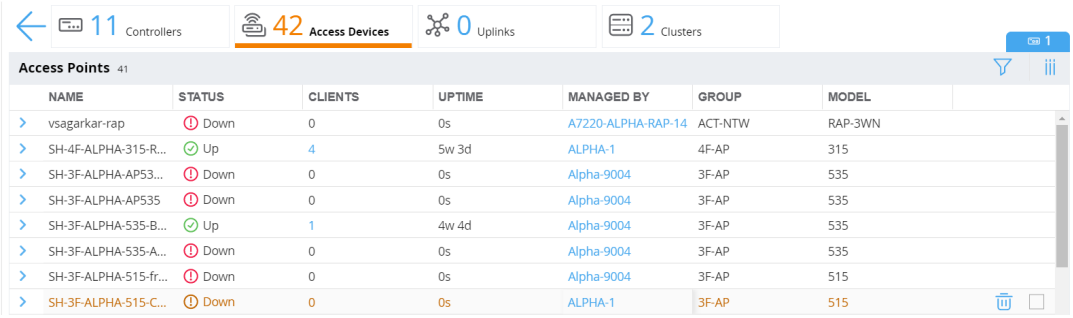
You can click on the **Controllers** tab to view the status of each controller connected to the cluster. It also displays information like the software version each controller is running and the number or APs and clients attached to each controller.

Figure 56 *Cluster Dashboard Controller Information*



Similarly, you can click on **Access Devices** to view the status of each AP connected to the cluster. It also displays information like Uptime, which group it belongs to, and the controller it is connected to.

Figure 57 *Cluster Dashboard Access Points Information*



You can also expand each controller or AP for more details on the controller or AP to view the name, IP addresses, MAC addresses, and so on.

Figure 58 *Controller Detailed View*

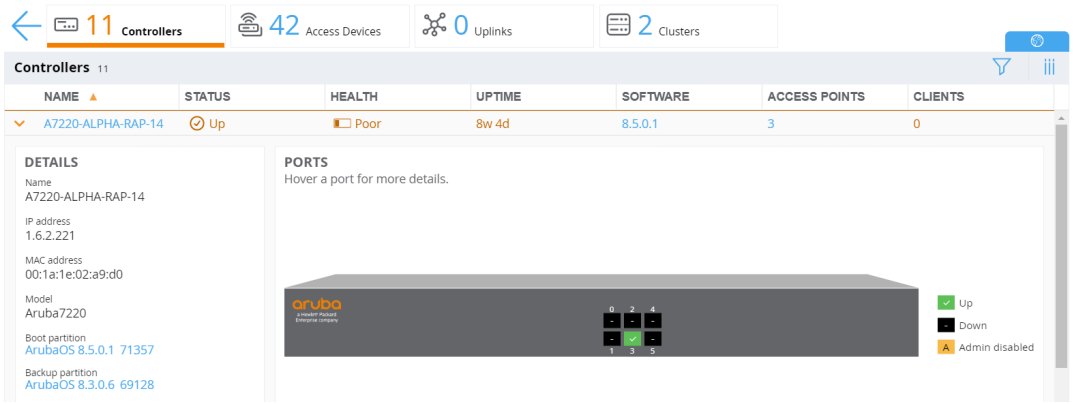
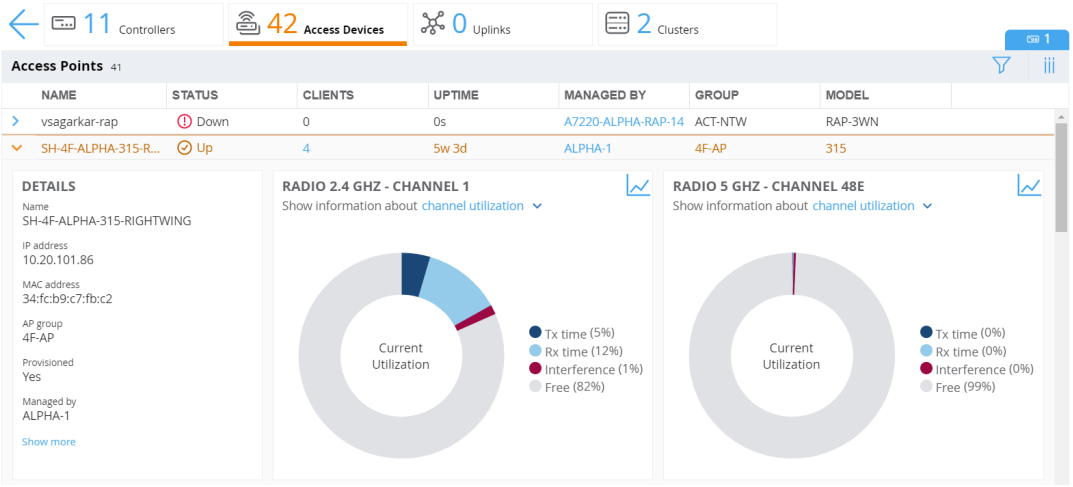


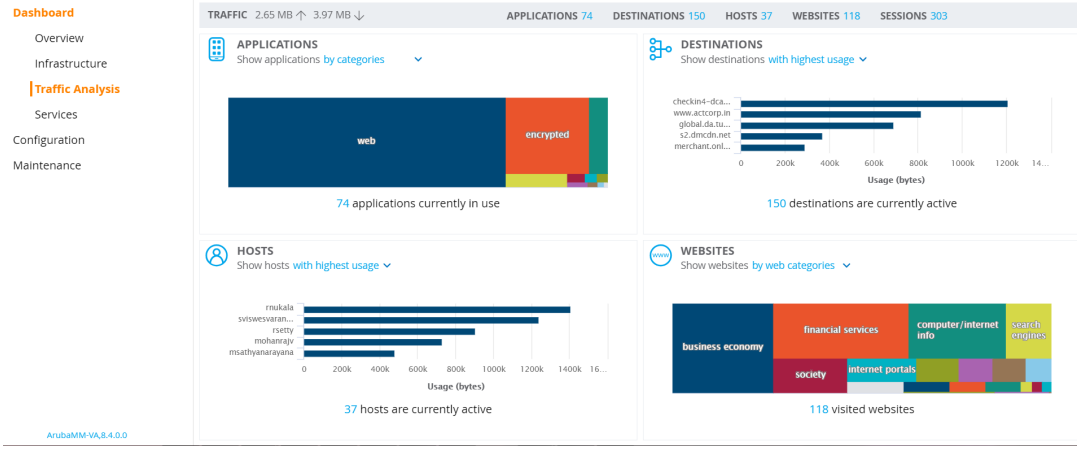
Figure 59 *Access Point Detailed View*



Traffic Analysis

The **Traffic Analysis** page provides the summary of APPLICATIONS, DESTINATIONS, HOSTS, and WEBSITES features. See Figure 60 for **Traffic Analysis** page.

Figure 60 *Traffic Analysis Page*



You can click the hyperlinked number of a particular feature to navigate to its table to view more information.

NOTE

The **Traffic Analysis** dashboard application visibility feature is supported only in 7000 Series, 7200 Series , and x86 managed devices, and requires WebCC and PEFNG license.

The **Traffic Analysis** dashboard contains the following windows:

- **APPLICATIONS**—This window displays the summary of all applications in the managed device. This is the default page. Click the Applications window or hyperlinked number to navigate to the **Applications** table and view the details of the applications currently in use. For more information, see Applications on page 432.

You can view the following information using the **Show applications** drop-down list.

- **By categories** — Displays all the applications by categories.
- **With highest usage** — Displays five applications with highest number of bytes transmitted or received, in decreasing order. Click the horizontal bar to display the details of the selected application.
- **With most hosts** — Displays five applications with highest number of hosts, in decreasing order. Click the horizontal bar to display the details of the selected application.

- **DESTINATIONS**—This window displays the summary of all active destinations in the managed device. This is the default page. Click the horizontal bar or hyperlinked number to navigate to the **Destinations** table and view the details of the active destinations. For more information, see Destinations on page 433.

You can view the following information using the **Show destinations** drop-down list.

- **With highest usage** — Displays five destinations with highest number of bytes transmitter or received, in decreasing order. Click the horizontal bar to display the details of the selected destination.
- **With most hosts** — Displays five destinations currently accessed by highest number of hosts, in decreasing order. Click the horizontal bar to display the details of the selected destination.

- **WEBSITES**—This window displays the summary of all websites visited using the managed device in the network hierarchy. This is the default page. For more information, see Websites on page 436.

You can view the following information using the **Show websites** drop-down list.

- **By reputation** — Displays percentage of traffic based on reputation or score of web traffic in the managed device. The reputation levels are Trustworthy, Low risk, Moderate risk, Suspicious, High risk, and Unknown. Click the pie chart to display details of the selected reputation level.
- **By web categories** — Displays the number of bytes transmitted or received by web categories in tree chart presentation. Click on the rectangle tile of a selected web category to show the number of bytes transferred for the selected web category that is grouped by reputation.
- **With highest usage** — Displays five websites with the highest number of bytes transmitted or received, in decreasing order in chart presentation. Click the horizontal bar to display details of the selected website.
- **With most hosts** — Displays five websites with the highest number of clients currently connected, in decreasing order in chart presentation. Click the horizontal bar to display details of the selected website.

- **HOSTS**—This window displays the summary of all active hosts in the managed device. This is the default page. Click the horizontal bar or hyperlinked number to navigate to the **Hosts** table and view the details of the active hosts. For more information, see Hosts on page 435.

You can view the following information using the **Show hosts** drop-down list.

- **With highest usage** — Displays five hosts with highest number of bytes transmitted or received, in decreasing order. Click the horizontal bar to display the details of the selected host.

- **By OS** — Displays the number of hosts that are running each type of OS. Click the OS area in the donut chart to display the hosts running the selected OS type.
- **By Role** — Displays the number of hosts with individual roles. Click the role area in the donut chart to display the hosts of the selected role.

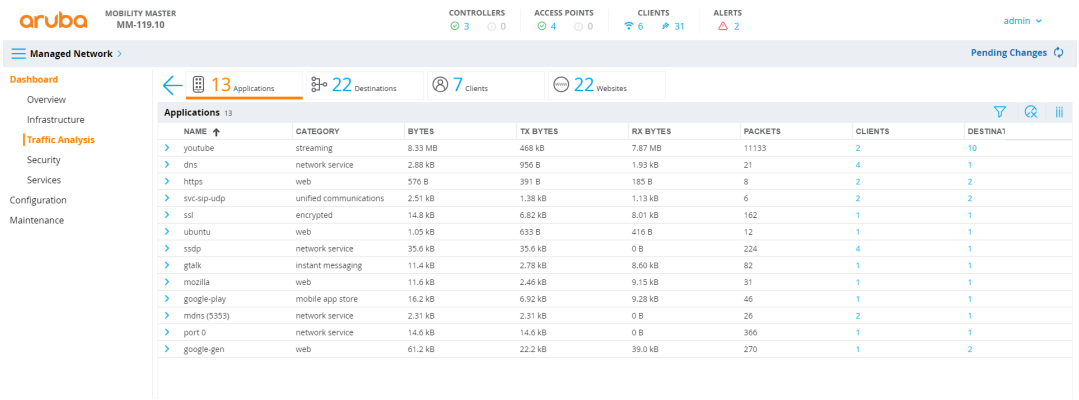
Applications

Applications performs DPI of local traffic and detects over 1500 applications on the network. **Applications** allows you to configure both application and application category policies within a given user role.

Enable DPI to enhance the benefit of the existing visualization or dashboard, To enable DPI, see the Enabling DPI section.

Navigate to **Dashboard > Traffic Analysis** and click **APPLICATIONS** icon. The **Applications** page displays the summary of all the applications in the managed device. See Figure 61 for **Applications** page.

Figure 61 *Applications Page*



Action Bar

The Action bar displays the total number of applications depending on filters applied. The action bar includes action icons namely, **Show/Hide table filters**, **Block/Unblock applications**, **Customize columns**, and **Set application QoS**.

You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the applications in the table that you want to view.
- **Block/Unblock applications** — This button allows you to permit or deny an application or an application category for a given role. You can create global and per-role rules. For example, you can block the YouTube application, which belongs to the Streaming application category for the guest role within the enterprise.

To **Block/Unblock applications**, perform the following steps:

1. Select a application from the **Applications** table.

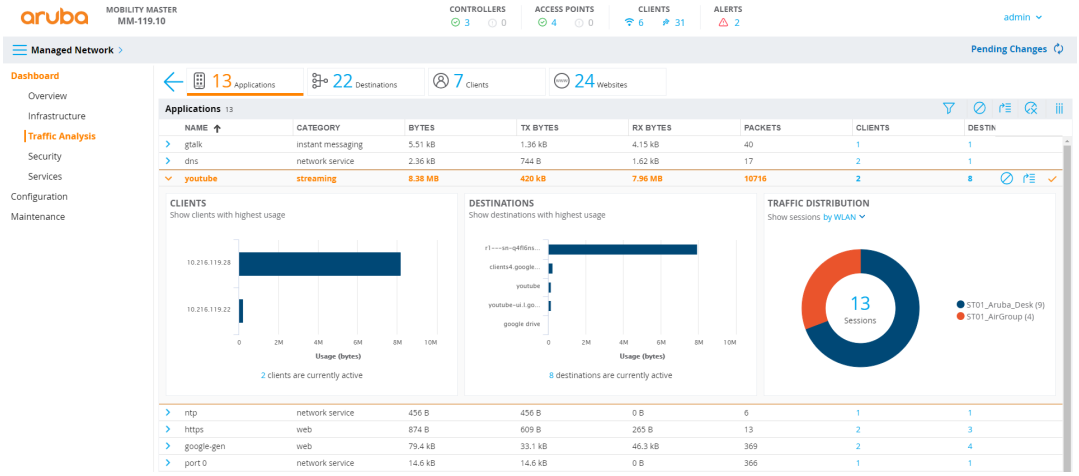
Details

Expand the application from the **Applications** table to view the detailed information of individual application. See Figure 62 for Details page.

The application table displays the following details:

- **HOSTS** — Displays five hosts with highest application usage, in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the **Hosts** table and view the details of the hosts using this application.
- **DESTINATIONS** — Displays five destinations with highest usage, in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the **Destinations** table and view the details of the destination using this application.
- **TRAFFIC DISTRIBUTION** — Displays the distribution of traffic. You can view the following information using the **Show sessions** drop-down list.
  - **By WLAN** — Displays the number of sessions established by the application for each WLAN in a donut chart. Click the donut chart area or hyperlinked number to navigate to the **Sessions** table and view the details of the sessions established by this application.
  - **By OS** — Displays the number of sessions established by the application for each OS in a donut chart. Click the donut chart area or hyperlinked number to navigate to the **Sessions** table and view the details of the sessions established by this application.
  - **By Role** — Displays the number of sessions established by the application for each role in a donut chart. Click the donut chart area or hyperlinked number to navigate to the **Sessions** table and view the details of the sessions established by this application.

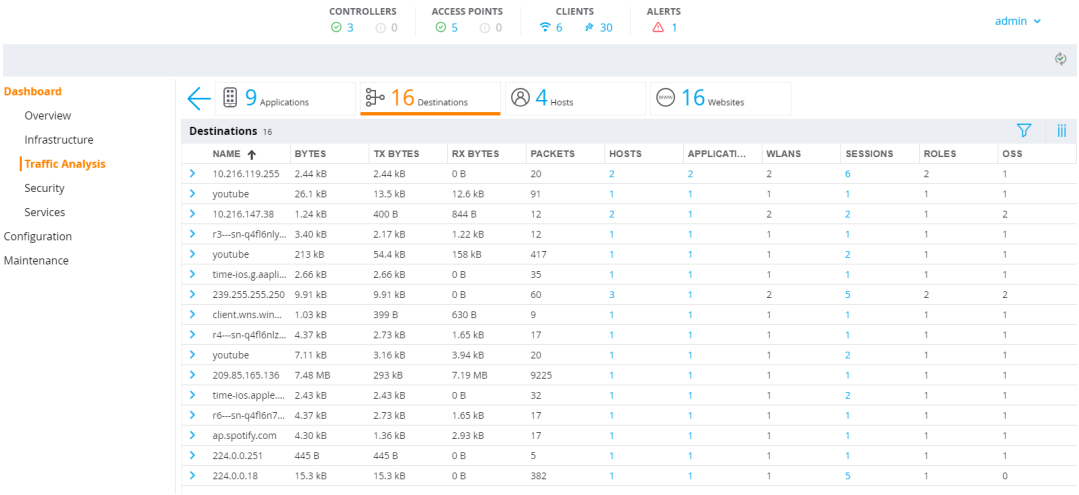
Figure 62 *Details Page*



Destinations

Navigate to **Dashboard > Traffic Analysis** and click **DESTINATIONS** icon. The **Destinations** page displays the summary of all the active destinations in the managed device. See Figure 63 for **Destinations** page.

Figure 63 *Destinations Page*



Action Bar

The Action bar displays the total number of destinations depending on the filters applied. The Action bar includes action buttons namely, **Show/Hide table filters**, and **Customize columns**.

You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the destinations in the table that you want to view.
- **Customize columns** — Click **Customize columns** icon to select the table columns that you want to view.

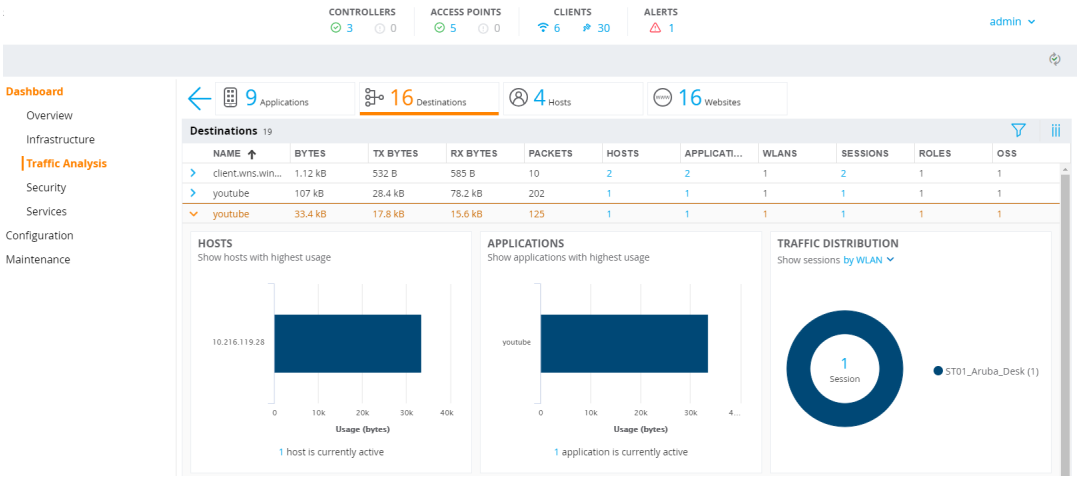
Details

Expand the destination from the **Destinations** table to view the detailed information of individual destination. See Figure 64 for Details page.

The destination table displays the following details:

- **HOSTS** — Displays five hosts with highest destination usage, in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the **Hosts** table and view the details of the hosts using this application.
- **APPLICATIONS** — Displays five applications with highest usage, in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the **Applications** table and view the details of the application used by the destination.
- **Traffic Distribution** — Displays the distribution of traffic. You can view the following information using the **Show sessions** drop-down list.
  - **By WLAN** — Displays the number of sessions established by the destination for each WLAN in a donut chart. Click the donut chart area or hyperlinked number to navigate to the **Sessions** table and view the details of the sessions established by this destination.
  - **By OS** — Displays the number of sessions established by the destination for each OS in a donut chart. Click the donut chart area or hyperlinked number to navigate to the **Sessions** table and view the details of the sessions established by this destination.
  - **By Role** — Displays the number of sessions established by the destination for each role in a donut chart. Click the donut chart area or hyperlinked number to navigate to the **Sessions** table and view the details of the sessions established by this destination.

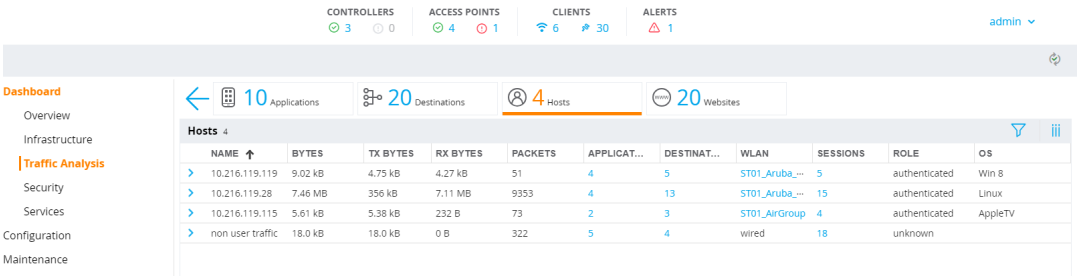
Figure 64 *Details Page*



Hosts

Navigate to **Dashboard > Traffic Analysis** and click **HOSTS** icon. The **Hosts** page displays the summary of all the active hosts in the managed device. See Figure 65 for **HOSTS** page.

Figure 65 *Hosts Page*



Action Bar

The Action bar displays the total number of hosts depending on the filters applied. The Action bar includes action buttons namely, **Show/Hide table filters**, and **Customize columns**.

You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the clients in the table that you want to view.
- **Customize columns** — Click **Customize columns** icon to select the table columns that you want to view.

Details

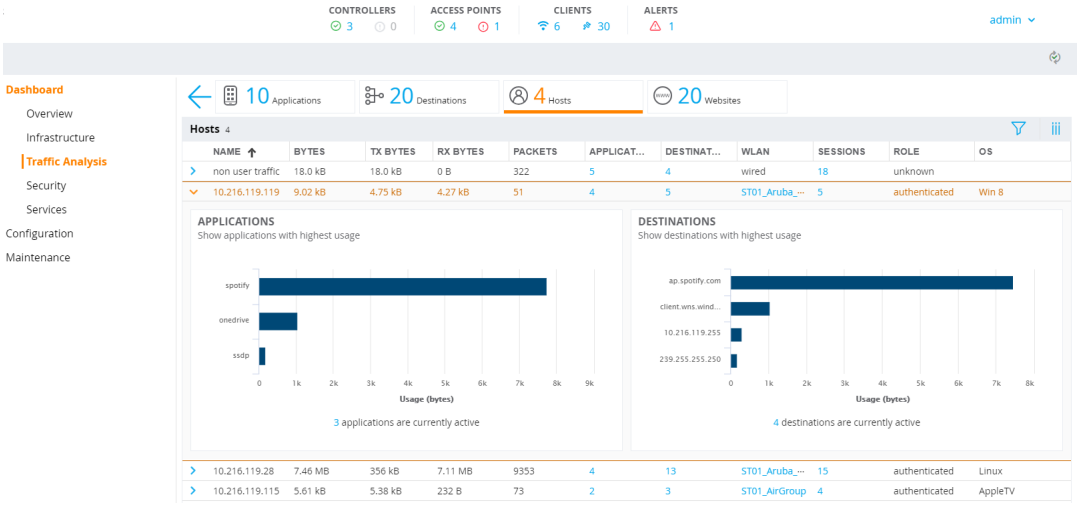
Expand the host from the **Hosts** table to view the detailed information of individual host. See Figure 66 for Details page.

The host table displays the following details:

- **APPLICATIONS** — Displays five applications with highest usage, in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the **Applications** table and view the details of the applications used by the selected host.

- **DESTINATIONS** — Displays five destinations with highest usage, in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the **Destinations** table and view the details of the destinations used by the selected host.

Figure 66 *Details Page*



Websites

Navigate to **Dashboard > Traffic Analysis** and click the **WEBSITES** tab. The implementation of **WebCC** feature can be viewed in this tab. **WebCC** uses a cloud-based service to dynamically determine the types of websites being visited, and their safety.

NOTE

The **WebCC** feature requires the WebCC subscription license.

When the **WebCC** feature is enabled, all web traffic (http and https) is classified. The classification is done in data path as the traffic flows through the managed device and updates dynamically.

Starting from ArubaOS 8.4.0.0, the **WebCC** feature supports classification of both IPv4 and IPv6 sessions.

Aruba has partnered with Webroot®, a Web classification service to provide the **WebCC** feature in the Mobility Master. Aruba uses the URL database of Webroot and the cloud look-up service to classify the web traffic. Aruba uses Webroot classified categories and score for web categories and reputation for **WebCC**. The following **Websites** Reputation table lists the risk level and score associated to each reputation level:

Table 73: *Websites Reputation Table*

Risk Level	Score
High-risk	1 - 20
Suspicious	21 - 40
Moderate-risk	41 - 60
Low-risk	61 - 80
Trustworthy	81 - 100

As indicated inTable 73, if the risk level of the content reputation goes high, the score goes down. So, when a WebCC reputation is configured to be denied in a policy, all traffic with the specified risk level and higher are denied. Similarly, when a WebCC reputation is configured to be permitted in a policy, all traffic with the specified risk level and lower are permitted. For example, if a policy is configured to deny moderate-risk traffic, then all the traffic categorized under moderate-risk, suspicious, and high-risk levels are denied. If a policy is configured to permit moderate-risk traffic, then all the traffic categorized under moderate-risk, low-risk, and trustworthy levels are permitted.

Action Bar

The Action bar displays the total number of websites depending on the filters applied. The Action bar includes action buttons namely, **Show/Hide table filters**, **Block/Unblock web category**, **Set web category QoS**, and **Customize columns**.

You can perform the following tasks on this page:

- **Show/Hide table filters** — Click the **Show/Hide table filters** icon to filter the websites by **NAME**, **CATEGORY**, or **REPUTATION**.
- **Block/Unblock web category** — This button allows you to permit or deny a website or web category for a given role.

WebCC Operation Modes

The WebCC can operate in one of the following modes:

- **Centralized**—In the Centralized mode, the Mobility Master downloads the URL entry database from Webroot®. This is the default mode.
- **Distributed**—In the distributed mode, the managed device downloads the URL entry database directly from Webroot®. This method is useful when the Mobility Master is unreachable by the managed device. You can configure the WebCC operation in **Distributed** mode using the CLI. For more information, see s.

The current policy enforcement model in Aruba relies on L3 or L4 information of the packet or L7 information with DPI support to apply rules. WebCC complements this as the user is allowed to apply firewall policies based on web content category and reputation.

Benefits of WebCC

The benefits of **WebCC** are as follows:

- Prevention of malicious malware, spyware, or adware by blocking known dangerous websites
- Visibility into web content category-level
- Visibility into web sites accessed by the user

Accessing Websites in Dashboard

Navigate to **Dashboard > Traffic Analysis** to view the **WEBSITES** tab.

The **Websites** tab includes the following containers:

- **By reputation:** The reputation pie chart shows the percentage of traffic based on reputation or score of web traffic in the managed device. The reputation levels are Trustworthy, Low-Risk, Moderate-Risk, Suspicious, and High-Risk. If there is no traffic on a specific reputation, then the corresponding reputation does not appear in the pie chart. The circles in this chart are click-able. Clicking on a circle filters rest the of the page data with the selected reputation.

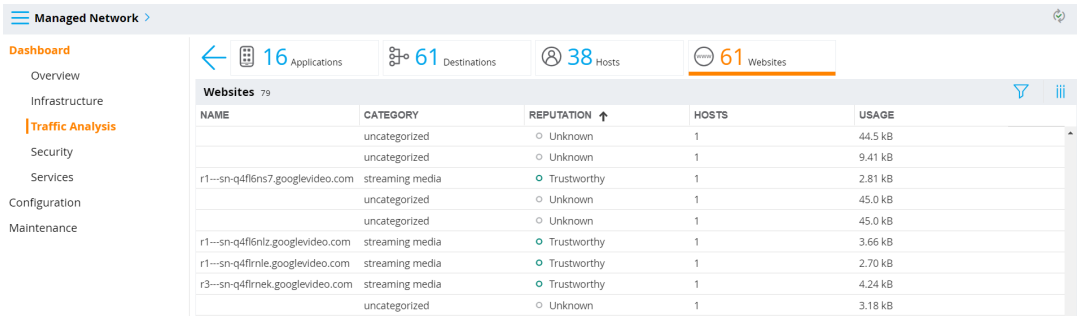
- **By web categories:** This chart shows traffic for web categories in tree chart presentation. All boxes in this chart are click-able. Clicking on a box filters rest of page data with the clicked web category as filter.
- **With highest usage**: When the **websites with highest usage** option is selected, the Websites chart displays five websites with the highest number of bytes transferred, in decreasing order:
- **With most hosts:** When the **websites with most hosts** option is selected, the Websites chart displays five websites with the highest number of clients currently connected, in decreasing order.

  The **Websites** table of the selected web category includes the following five columns:

  - **NAME**: Lists the website.
  - **CATEGORY**: Lists the type of web category for a particular website.
  - **REPUTATION**: Lists the reputation of web traffic.
  - **HOSTS**: Lists the number of clients currently connected for a website.
  - **USAGE**: Lists the number of bytes transferred by a website.

The following figure shows the **Websites** table:

Figure 67 *Websites Table*



WebCC Configuration in the WebUI

The following topics provide information on enabling WebCC, configuring new policies, and configuring WebCC bandwidth contract using the WebUI.

Enabling WebCC

You can enable WebCC using the **Applications** or **Global Settings** page.

The following procedure describes how to enable WebCC using the **Applications** page:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Roles & Policies > Applications** page.
2. In the **Application Visibility** accordion, select the **Web content classification** check box to enable WebCC.

The following procedure describes how to enable WebCC using the **Global Settings** page:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Services > Firewall >Global Settings**.
2. In the **Global Settings** accordion, select the **Enable Web Content Classification** check box.
3. On enabling, WebCC, the WebCC usage information is sent to Aruba at every 7 days interval.
4. From the **Connect to classification server using** drop-down list, select **IPv4** or **IPv6** based on your preference.
5. Click **Submit**.

6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.



When you enable WebCC from the **Applications** page, the **Enable Web Content Classification** option in the **Global Settings** page is enabled automatically along with the default value of **IPv4**. You can select **IPv6** from the **Connect to classification server using** drop-down list if you want to provide IPv6 support for WebCC.

New Policy Configuration

The following procedure describes how to configure a new policy and create an ACL rule with web category and reputation:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Roles & Policies > Policies** page.
2. In the **Policies** table, click **+** to open the **Add Policy** window .
3. In the **Add Policy** window, select **Session** from the **Policy type** drop-down list and enter a name for **Policy name**.
4. Click **Submit**.
5. Select the newly added policy from the **Policies** table to display the **Policies > <policy name>** section below.
6. In the **Policies > <policy name>** table, click **+** to open a the **New rule for <policy name>** window.
7. For **Rule Type**, select **Application** and click **OK**.
8. In the **Roles > <policy name> > New application Rule** section:
   a. From the **IP version** drop-down list, select **Any**, **IPv4**, or **IPv6** based on your preference.
   b. From the **Source** drop-down list, select **Any** or **User**.
   c. From the **Destination** drop-down list, select **Any** or **User**.
   d. From the **Scope** drop-down list, select **Web Category/Reputation**.



Based on your selection in the **Scope** drop-down list, the **Web category** or **Web reputation** parameter is displayed.

   e. From the **Web category** or **Web reputation** drop-down list, select the appropriate value.
   f. From the **Action** drop-down list, select **Deny** to not allow users to access this web category or **Permit** to allow users to access the web category.
   g. For **TOS**, enter a value.
   h. From the **Time range** drop-down list, select a suitable time range for which you want the policy to be active or valid. Alternatively, you can also create a new time range but clicking + in this field
   i. From the **802.1p priority** drop-down list, select a priority from 1 to 7.
   j. For **Options**, select **Log, Mirror**, and **Blacklist**, or any of these options that is applicable.
9. Click **Submit**.

10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

WebCC Configuration in the CLI

The following topics provide information on enabling WebCC, configuring new policies, and configuring WebCC bandwidth contract using the CLI.

Enabling WebCC

Use the following command to enable WebCC:

```
(host) [mynode] (config) #firewall
(host) [mynode] (config-submode)#web-cc
```

Use the following command to view the WebCC configuration:

```
(host) [md] #show firewall
Global firewall policies
------------------------
Policy                                        Action          Rate       Port
------                                        ------          ----       ----
Enforce TCP handshake before allowing data    Disabled
Prohibit RST replay attack                    Disabled
.....
Web Content Classification                    Enabled
....
```

Use the following command to configure WebCC per-role:

```
(host) [mynode](config-submode) #web-cc
```

Use the following command to configure the WebCC operation mode to **distributed** from the default **centralized** mode:

```
(host) [mynode] (config) #webcc distributed
```

New Policy Configuration

The new CLI extends the existing policy configuration to take web category, reputation, or both.

Use the following command to configure a new policy to create ACL rule with web category and reputation:

```
(host) (mynode)(config-submode) #source destination proto-port/service/app/app-group
<name> webcc-category <ctgry> webcc-reputation <score> action [log | mirror | time-range]
```

Starting from ArubaOS 8.4.0.0, for IPv6 sessions, use the following command to configure a new policy to create ACL rule with web category and reputation:

```
(host) (mynode)(config-submode) #ipv6 source destination proto-port/service/app/app-group
<name> web-cc-category <ctgry> web-cc-reputation <score> action [log | mirror | time-
range
```

The following actions are supported when web category/reputation is selected:

- Deny
- Permit
- Blacklist
- Classify-media
- Disable-scanning
- Dot1q-priority
- Log
- Mirror
- Queue

- Time-range
- TOS

Example for WebCC policy configuration is as follows:

```
(host) [mynode] (config) #ip access-list session url-filter
(host) [mynode] (config-submode)#any any web-cc-category translation permit
(host) [mynode] (config-submode)#any any web-cc-reputation high-risk deny
(host) [mynode] (config-submode)#any any any deny
```

Example for WebCC policy configuration only for **http** traffic running on TCP 80, the above ACL is modified as follows in datapath for pre-classification ACL scan:

```
(host) [mynode] (config) #ip access-list session url-filter
(host) [mynode] (config-submode)#any any tcp {80} permit
(host) [mynode] (config-submode)#any any tcp {80} deny
(host) [mynode] (config-submode)#any any any deny
```

Post-classification, ACL look-up will have the ACL as follows:

```
(host) [mynode] (config) #ip access-list session url-filter
(host) [mynode] (config-submode)#any any tcp {80} WebCCCtgID 40 WebCCRep 1-100 permit
(host) [mynode] (config-submode)#any any tcp {80} WebCCRep 1-20 deny
(host) [mynode] (config-submode)#any any any deny
```

In case there exists an ACL rule to deny or permit a specific web category but is required to make an exception to allow or deny a specific URL or website, then this can be accomplished by configuring in the following manner:

1. First define a netdestination with one or more URLs to whitelist or blacklist
   ```
   (host) [mynode] (config) #netdestination search
   (host) [mynode] (config-submode) #name www.google.com
   (host) [mynode] (config-submode) #name www.bing.com
   (host) [mynode] (config-submode) #exit
   ```
2. Apply this netdestination to an ACL
   ```
   (host) [mynode] (config) #ip access-list session whitelist
   (host) [mynode] (config-submode)#any alias search tcp 80 permit
   (host) [mynode] (config-submode)#any alias search tcp 443 permit
   ```
3. Apply this ACL to a user-role. The position of this ACL should be at the top. However, with global or role-specific default ACLs this wouldn't be possible.
   ```
   (host) [mynode] (config) #user-role guest2
   (host) [mynode] (config-submode) #access-list session whitelist
   ```



If there a WebCC or app rule that is applicable globally across user-roles then you cannot override such behavior. This is a limitation.

WebCC Bandwidth Contract Configuration

With this feature, ArubaOS supports configuring WebCC category and reputation based bandwidth contract configuration or enforcement. This can be enforced globally for all user-roles, or can be enforced per user-role.

Use the following command to apply global WebCC based bandwidth contracts using the CLI:

```
(host) (mynode) (config) #web-cc global-bandwidth-contract {webcc-category|webcc-
reputation}{upstream|downstream}{kbits <value>|mbits <value>}
```

Use the following command to apply AAA bandwidth contracts using the CLI:

```
(host) (mynode) (config) #aaa bandwidth-contract webcc mbits <value>
```

Use the following command to apply role-specific web-cc based bandwidth contracts using the CLI:

```
(host) (mynode) (config) #user-role webcc
```

```
(host) (mynode) (config-role) #bw-contract {webcc-category|webcc-reputation}<name>
<contract> {upstream|downstream}{kbits <value>|mbits <value>}
```

Debugging

The following **show** commands are introduced as part of this feature:

- **show web-cc category all**: Displays all WebCC categories
- **show web-cc reputation:** Displays WebCC reputation
- **show web-cc stats:** Displays the statistics of WebCC
- **show web-cc status:** Display the status of Web-CC
- **show web-cc global-bandwidth-contract:** Displays configured WebCC bandwidth contract
- **show datapath web-cc:** Displays md5, web category, reputation, and age for each URL
- **show datapath web-cc counters:** Displays the number of URLs in cache, Classified and Unclassified sessions.
- **show datapath session web-cc:** Displays Internal Flags, Pre Classification ACE Index, and Post Classification ACE Index
- **show datapath session ipv6 web-cc counters:** Displays IPv6 datapath session statistics of WebCC.
- **show gsm debug channel web_cc_info**: Lists md5, Category, and Reputation for each URL. GSM entries are populated as and when URL cache entry is learned, and it is used for reporting the actual URLs being associated with user session entries.

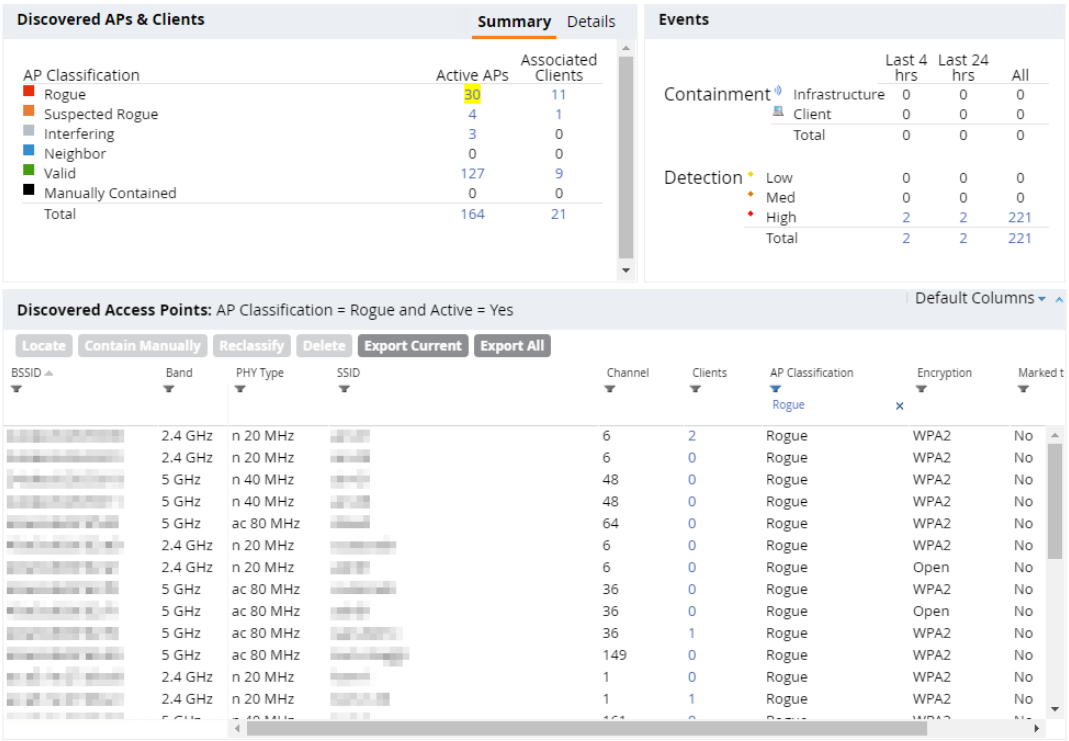
The following **clear** commands are introduced as part of this feature:

- **clear web-cc cache <md5_1> <md5_2>** : Clears the WebCC cache entry from both data plane and GSM.
- **clear web-cc stats**: Clears all WebCC statistics.
- **clear datapath web-cc counters**: Clears configuration values and statistics in the WebCC datapath module.

Security

The **Security** page displays the summary of detected radios, detected clients, events , and blacklisted clients in your network. See Figure 68 for **Security** Page.

Figure 68 *Security Page*



The **Security** dashboard contains the following windows:

- **Detected Radios** — Displays the active detected radios in the managed device. This is the default page. The detected radios is categorized as: **Authorized**, **Neighbor**, **Interfering**, **Suspected Rogue**, **Rogue**, and **Contained**. Click the donut chart area or hyperlinked number to navigate to the **Detected Radios** table and view the details of the detected radios in the managed device. For more information, see Detected Radios on page 444.
- **Detected Clients** — Displays the active detected clients in the managed device. This is the default page. The detected clients is categorized as: **Authorized**, **Interfering**, and **Contained**. Click the donut chart area or hyperlinked number to navigate to the **Detected clients** table and view the details of the detected clients in the managed device. For more information, see Detected Clients on page 445
- **Events** — Displays the information about all detected and containment events that occur during specified time frame. For more information, see Events on page 446

You can view the following information using the **Grouped by** drop-down list.

- **Severity** — Displays low, medium and high severity of the events occurred in **last 4 hours**, **last 24 hours**, or **anytime**. This is the default page. Click the vertical bar or hyperlinked number to navigate to the **Events** table and view the details of the events you wish to view.
- **Target type** — Displays the details of the events by infrastructure and client occurred in **last 4 hours**, **last 24 hours**, or **anytime**. Click the vertical bar or hyperlinked number to navigate to the **Events** table and view the details of the events you wish to view.
- **Feature type** — Displays the details of the events by detection and containment occurred in **last 4 hours**, **last 24 hours**, or **anytime**. Click the vertical bar or hyperlinked number to navigate to the **Events** table and view the details of the events you wish to view.
- **Blacklisted Client** — Displays the clients that are blacklisted in stand-alone controllers or in Mobility Master-Managed Device topology. This is the default page. Click the **Blacklist** icon or donut chart

area or hyperlinked number to navigate to the **Blacklisted Clients** table and view the details of the clients that are blacklisted in the managed device.

Detected Radios

Navigate to **Dashboard > Security** and click **DETECTED RADIOS** icon. The **Detected Radios** page displays the summary of all the detected radios in the managed device. See Figure 69 for **Detected Radios** page.

Figure 69 *Detected Radios*

Action Bar

The Action bar displays the total number of detected clients depending on the filters applied. The Action bar includes action buttons namely, **Show/Hide table filters**, **Reclassify detected radios**, and **Customize columns**.

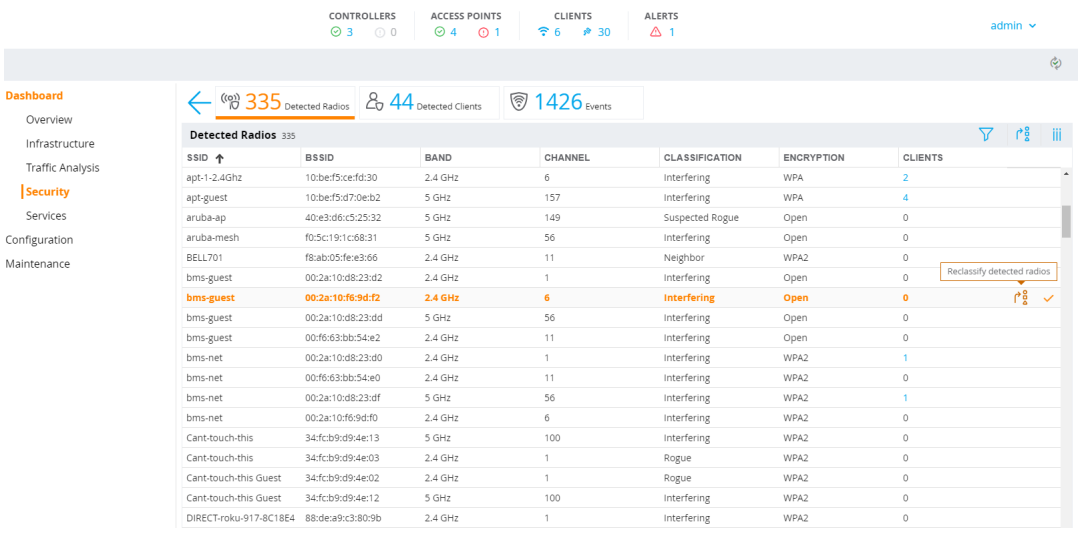
You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the detected clients in the table that you want to view.
- **Reclassify detected radios** — Allows you to reclassify a single detected radio or all the selected detected radios.

To **Reclassify detected radios**, perform the following steps:

1. Select a detected radio from the **Detected Radios** table.
2. Click **Reclassify detected radios** icon as shown in Figure 70.

Figure 70 *Reclassify Detected Radios*



SSID	BSSID	BAND	CHANNEL	CLASSIFICATION	ENCRYPTION	CLIENTS
apt-1-2.4Ghz	10:be:f5:ce:fd:30	2.4 GHz	6	Interfering	WPA	2
apt-guest	10:be:f5:d7:0e:b2	5 GHz	157	Interfering	WPA	4
aruba-ap	40:e3:d6:c5:25:32	5 GHz	149	Suspected Rogue	Open	0
aruba-mesh	f0:5c:19:1c:68:31	5 GHz	56	Interfering	Open	0
BELL701	f8:ab:05:fe:e3:66	2.4 GHz	11	Neighbor	WPA2	0
bms-guest	00:2a:10:d8:23:d2	2.4 GHz	1	Interfering	Open	0
bms-guest	00:2a:10:f6:9d:f2	2.4 GHz	6	Interfering	Open	0
bms-guest	00:2a:10:d8:23:dd	5 GHz	56	Interfering	Open	0
bms-guest	00:f6:63:bb:54:e2	2.4 GHz	11	Interfering	Open	0
bms-net	00:2a:10:d8:23:d0	2.4 GHz	1	Interfering	WPA2	1
bms-net	00:f6:63:bb:54:e0	2.4 GHz	11	Interfering	WPA2	0
bms-net	00:2a:10:d8:23:df	5 GHz	56	Interfering	WPA2	1
bms-net	00:2a:10:f6:9d:f0	2.4 GHz	6	Interfering	WPA2	0
Cant-touch-this	34:fc:b9:d9:4e:13	5 GHz	100	Interfering	WPA2	0
Cant-touch-this	34:fc:b9:d9:4e:03	2.4 GHz	1	Rogue	WPA2	0
Cant-touch-this Guest	34:fc:b9:d9:4e:02	2.4 GHz	1	Rogue	WPA2	0
Cant-touch-this Guest	34:fc:b9:d9:4e:12	5 GHz	100	Interfering	WPA2	0
DIRECT-roku-917-8C18E4	88:de:a9:c3:80:9b	2.4 GHz	1	Interfering	WPA2	0

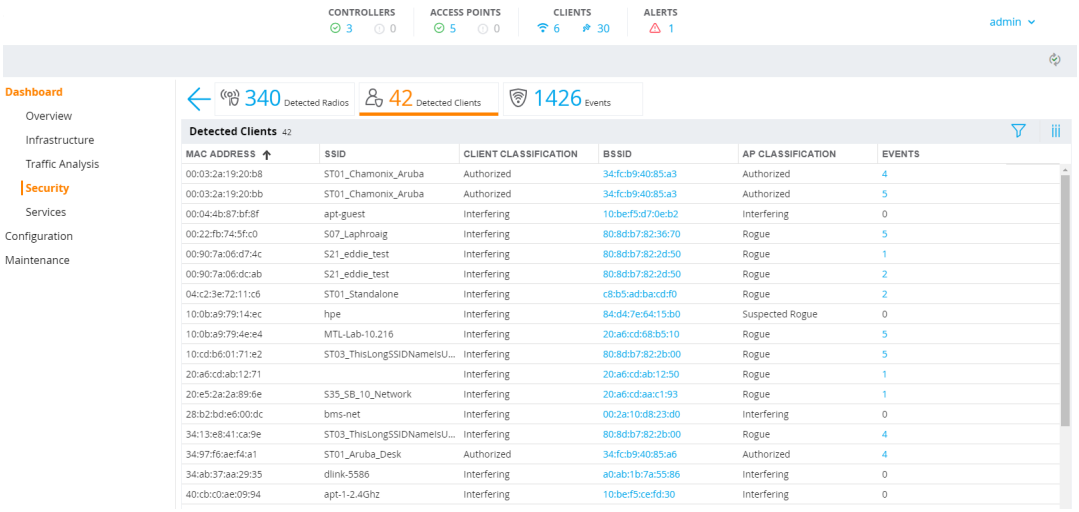
3. In the **Reclassify Detected Radios** dialog box, select classification from **Classification** drop-down list and click **Reclassify**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

- **Customize columns** — Click **Customize columns** icon to select the table columns that you want to view.

Detected Clients

Navigate to **Dashboard > Security** and click **DETECTED CLIENTS** icon. The **Detected Clients** page displays the summary of all the detected clients in the managed device. See Figure 71 for **Detected Radios** page.

Figure 71 *Detected Clients Page*



Action Bar

The Action bar displays the total number of detected clients depending on the filters applied. The Action bar includes action buttons namely, **Show/Hide table filters**, **Reclassify detected clients**, and **Customize columns**.

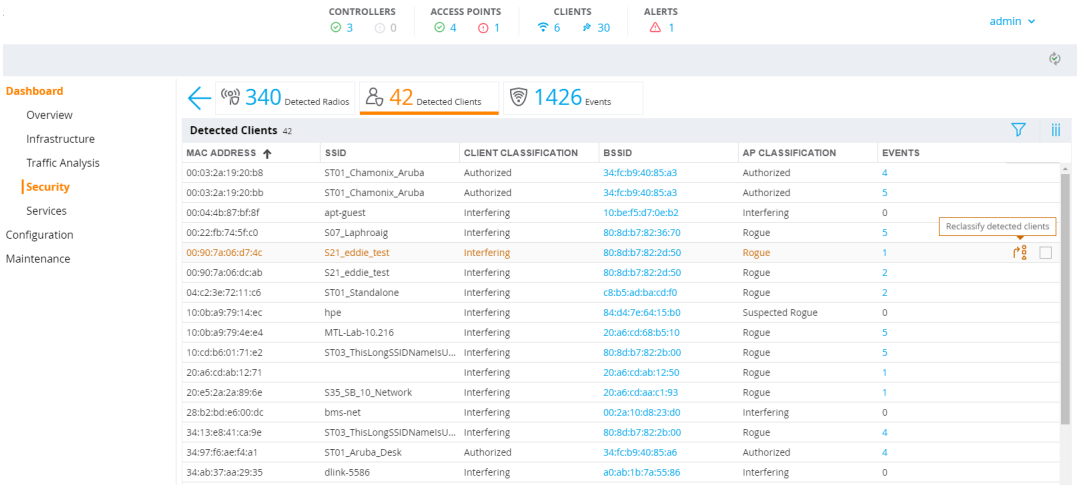
You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the detected clients in the table that you want to view.
- **Reclassify detected clients** — Allows you to reclassify a single detected client or all the selected detected clients.

To **Reclassify detected clients**, perform the following steps:

1. Select a detected client from the **Detected Clients** table.
2. Click **Reclassify detected clients** icon as shown in Figure 72.

Figure 72 *Reclassify Detected Clients*



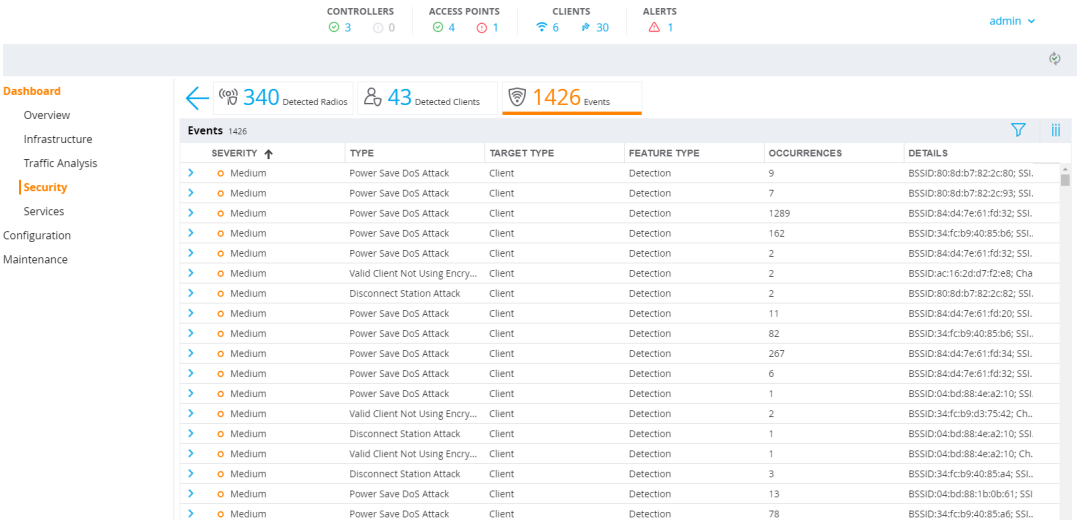
3. In the **Reclassify Detected Clients** dialog box, select classification from **Classification** drop-down list. Click **Reclassify**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

- **Customize columns** — Click **Customize columns** icon to select the table columns that you want to view.

Events

Navigate to **Dashboard > Security** and click **EVENTS** icon. The **Events** page displays the summary of all the events in the managed device. See for **Events** page.

Figure 73 *Events Page*



Action Bar

The Action bar displays the total number of wireless clients depending on the filters applied. The Action bar includes Action buttons namely, **Show/Hide table filters**, **Delete Event**, and **Customize columns**.

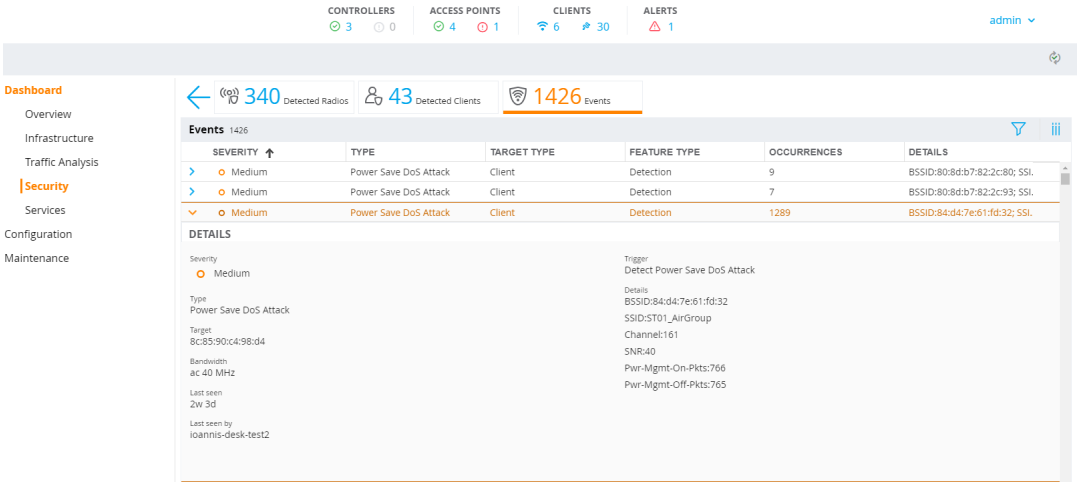
You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the wireless clients in the table that you want to view.
- **Delete Event** — Allows you to delete a selected event from the **Events** table.
- **Customize columns** — Click **Customize columns** icon to select the table columns that you want to view.

Details

Expand a event from the **Events** table to view the detailed information of individual event. See Figure 74 for Details page.

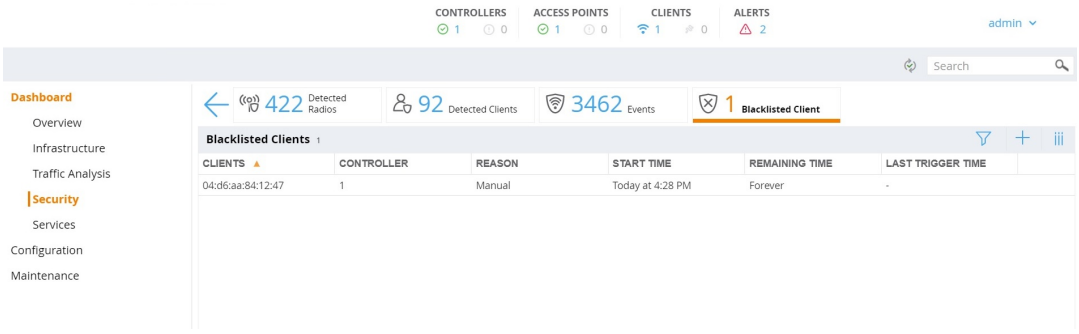
Figure 74 *Details Page*



Blacklisted Clients

Navigate to **Dashboard > Security** and click the **BLACKLIST** icon. The **Blacklisted Clients** page displays the summary of all the blacklisted clients in the managed device. See Figure 75 for **Blacklisted Clients** page.

Figure 75 *Blacklisted Clients*



Action Bar

The Action bar displays the total number of blacklisted clients depending on the filters applied. The Action bar includes action buttons namely, **Show/Hide table filters**, **Add to blacklist**, and **Customize columns**.

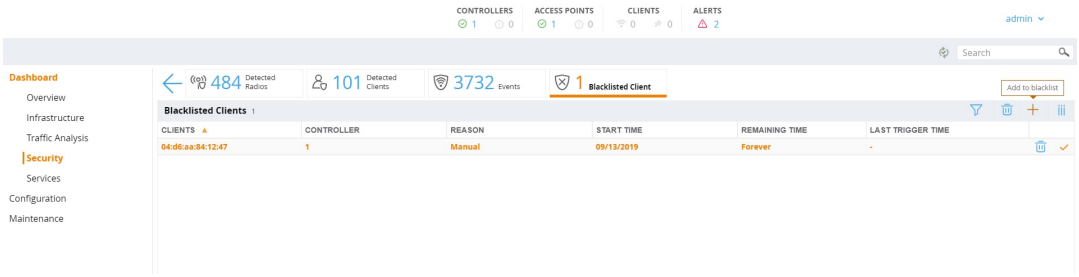
You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the blacklisted clients in the table that you want to view.
- **Add to blacklist** — Click **Add to blacklist** icon to add a client to the blacklist table.

To blacklist a client manually, perform the following steps:

1. Select a client from the **Wireless Clients** table.
2. Click **Add to blacklist** icon as shown in Figure 76.

Figure 76 *Add to Blacklist*



The **Add to Blacklist** pop-up window is displayed.

3. In the **Add to Blacklist** pop-up window, enter the MAC address of the client,
4. Click **Add**.

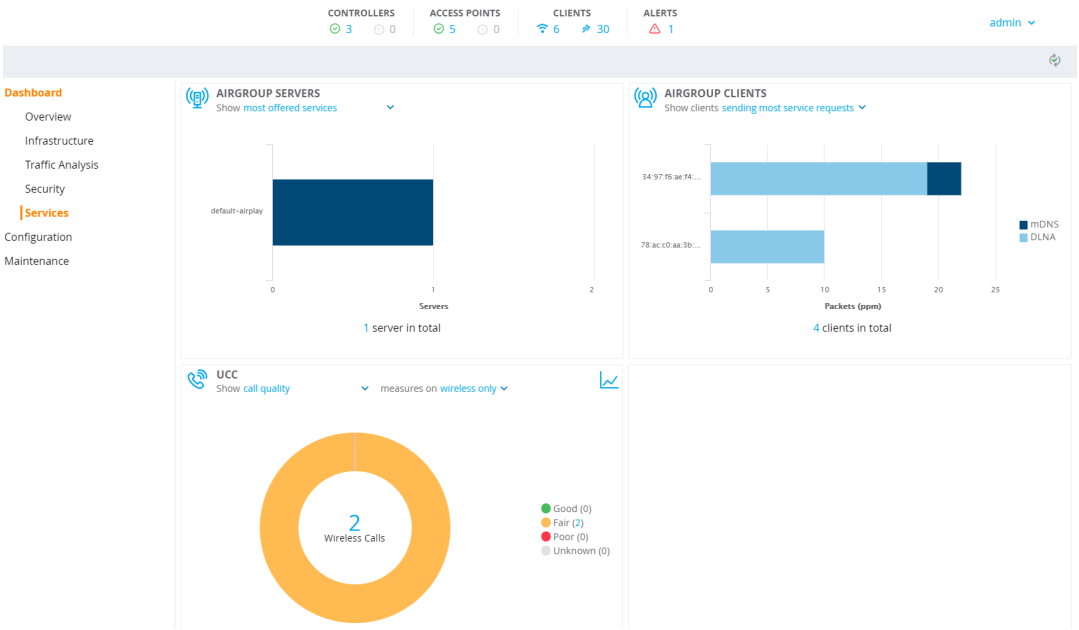
   The client is blacklisted and is listed in the **Blacklisted Clients** table.

- **Customize columns** — Click **Customize columns** icon to select the table columns that you want to view.

Services

The **Services** provides the summary of AirGroup servers, AirGroup clients, and UCC. See Figure 77 for **Services** page.

Figure 77 *Services Page*



The **Services** dashboard contains the following windows:

- **AIRGROUP SERVERS**—This window displays the summary of all the AirGroup servers in the network. This is the default page. For more information, see AirGroup Servers on page 450.

  You can view the following information using the **Show** drop-down list.

  - **Most offered services** — Displays five services advertised by the servers, in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the **Airgroup Servers** table and view the details of the AirGroup servers that offer the selected services.
  - **Servers with highest throughput** — Displays five AirGroup servers with highest throughput (bps), in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the AirGroup server table and view the details of the servers currently advertising AirGroup services.

  The **Servers with highest throughput** option from the **Show** drop-down list is displayed only if an AppRF license is installed and DPI is enabled.

- **AIRGROUP CLIENTS**—This window displays the summary of all the AirGroup clients in the network. This is the default page. For more information, see AirGroup Clients on page 450.

  You can view the following information using the **Show clients** drop-down list.

  - **Sending most service requests** — Displays five AirGroup clients that sends the most mDSN and DLNA control packets, in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the AirGroup clients table and view the details of the mDSN and DLNA control packets.
  - **With highest throughput** — Displays five AirGroup servers with highest throughput (bps), in decreasing order. Click the horizontal bar or hyperlinked number to navigate to the AirGroup clients table and view the details of the selected clients.

  The **With highest throughput** option from the **Show** drop-down list is displayed only if an AppRF license is installed and DPI is enabled.

- **UCC**— This window displays the summary of all the wireless calls made in the managed devices. The number of external calls made is displayed in the bottom right corner of the **UCC** window. This is the default page. For more information, see UCC on page 451.

  You can view the following information using the **Show** and drop-down list.

  - **Call Quality** — Displays the call quality of the wireless calls. The call quality is categorized as Good, Fair, Poor, or Unknown. Select **wireless only** or **end-to-end** from **measures on** drop-down list to display the call quality of wireless only or end-to-end calls in the managed device. Click the donut chart area or hyperlinked number to navigate to the **Wireless clients** table and view the details of the wireless only or end-to-end calls in the managed device.

  Click **Historical** icon in the top right corner of the window to display the number of active calls in the last 15 minutes.

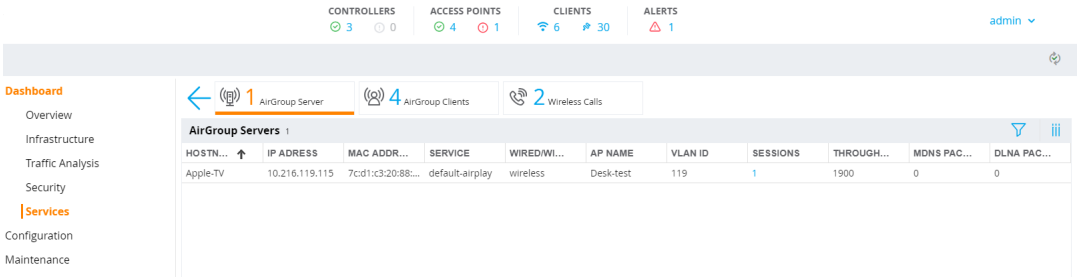
  - **Call Quality vs Client Health** Displays the co-relation between the VoIP call quality and the VoIP client health of every Unified Communication and Collaboration (UCC) call. Select **wireless only** or **end-to-end** from **measures on** drop-down list to display the co-relation between the VoIP call quality and the VoIP client health for wireless only or end-to-end calls. Click the dot on the plot chart to navigate to the **Wireless clients** table and view the details of the wireless only or end-to-end calls in the managed device.
  - **AP with most client** — Displays five access point groups with most number of calls. Click the horizontal bar to navigate to the **Wireless Calls** table and view the details of the selected wireless call.

- **OS** — Displays the number of clients that are running each type of operating systems. Click the OS area in the donut chart or hyperlinked number to navigate to the **Wireless calls** table and view the details of the wireless calls and the operating system type.
- **Application Protocol** — Displays the number of wireless calls and the type of application protocol. Click the application protocol area in the donut chart or hyperlinked number to navigate to the **Wireless calls** table and view the details of the wireless calls and the application protocol type.

AirGroup Servers

Navigate to **Dashboard > Services** and click **AIRGROUP SERVERS** icon. The **AirGroup Servers** page displays the summary of all the active AirGroup servers in the managed device. See Figure 78 for **AirGroup Servers** page.

Figure 78 *AirGroup Servers Page*



Action Bar

The Action bar displays the total number of AirGroup servers depending on the filters applied. The Action bar includes action buttons namely, **Show/Hide table filters**, and **Customize columns**.

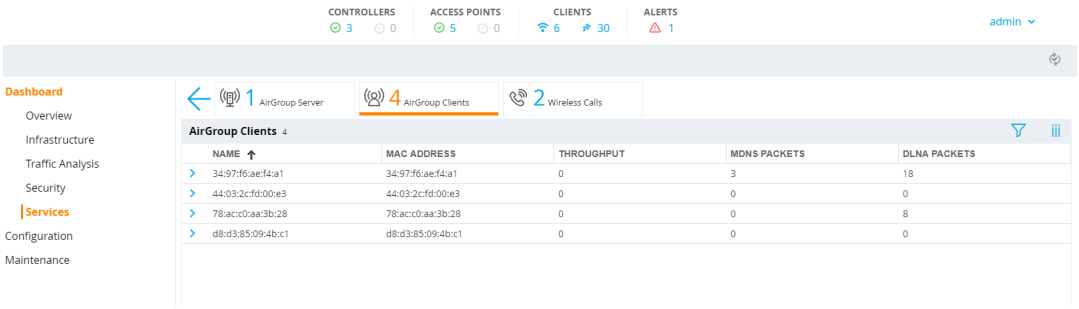
You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the AirGroup servers in the table that you want to view.
- **Customize columns** — Click **Customize columns** icon to select the table columns that you want to view.

AirGroup Clients

Navigate to **Dashboard > Services** and click **AIRGROUP CLIENTS** icon. The **AirGroup Clients** page displays the summary of all the active AirGroup clients in the managed device.See Figure 79 for **AirGroup Clients** page.

Figure 79 *AirGroup Clients Page*



Action Bar

The Action bar displays the total number of AirGroup clients depending on the filters applied. The Action bar includes action buttons namely, **Show/Hide table filters**, and **Customize columns**.

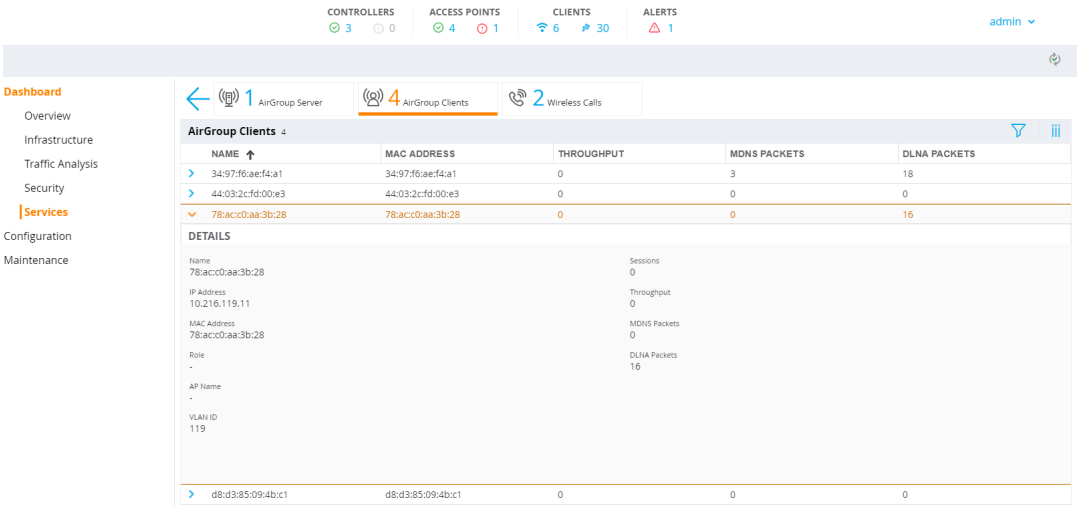
You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the AirGroup clients in the table that you want to view.
- **Customize columns** — Click **Customize columns** icon to select the table columns that you want to view.

Details

Expand a AirGroup client from the **AirGroup Clients** table to view the detailed information of individual AirGroup client. See Figure 80 for Details page.

Figure 80 *Details Page*



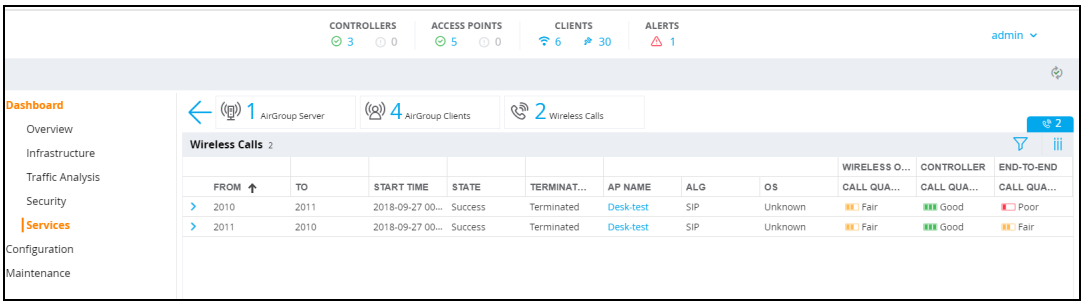
UCC

Navigate to **Dashboard > Services** and click **UCC** icon. The **Wireless Calls** page displays the summary of all the wireless calls made in the managed device. See for **Wireless Calls** page

NOTE

The UCC feature requires the PEFNG license.

Figure 81 *Wireless Calls Page*



Action Bar

The Action bar displays the total number of wireless calls depending on the filters applied. The Action bar includes action buttons namely, **Show/Hide table filters**, and **Customize columns**.

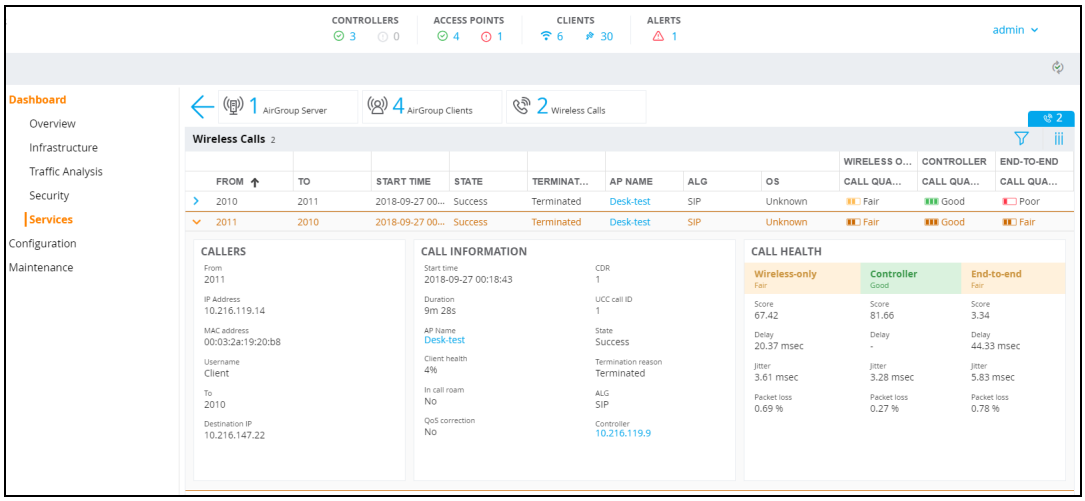
You can perform the following tasks on this page:

- **Show/Hide table filters** — Click **Show/Hide table filters** icon to filter and list the wireless clients in the table that you want to view.
- **Customize columns** — Click **Customize columns** icon to select the table columns that you want to view.

Details

Expand a wireless call from the **Wireless calls** table to view the detailed information of individual call like Callers, Call Information, and Call Health. See Figure 82 for Details page.

Figure 82 *Details Page*



The MultiZone feature allows organizations to have multiple and separate secure networks while using the same access point. It also allows AP to terminate to multiple managed devices that reside in different zones. A zone is a collection of managed devices under a single administration domain. The zone can have a single managed device or a cluster setup.

Traditionally, one AP was managed by a single zone where the configuration was generated on a master controller and synchronized across all other local controllers. Starting from ArubaOS 8.0.0.0, MultiZone AP is supported and an AP can be managed by multiple zones. Different zones can have different configurations. The managed devices in different zones do not communicate with one another.

Initially, when the AP is booted up, the first zone it contacts is called the Primary Zone. When the AP boots up on a managed device, and the primary zone managed device configures the AP including the BSS, radio channel, radio power, and other features. The primary zone can configure MultiZone profiles to enable the MultiZone feature.

Data zone is the secondary zone that an AP connects to after receiving the MultiZone configuration from the primary zone. If there are MultiZone profiles configured and associated in the AP group or AP name profile of the primary zone, then the AP enters MultiZone state and starts connecting with the specified data zones. Only one MultiZone profile per ap-group or ap-name can be attached. The data zone managed device must be configured with the same AP group or AP name profile as the primary zone. When the AP connects to the data zone managed devices, there is a flag in the HELLO message indicating that the AP is connecting to the zone as a data zone. The data zone managed device then can configure additional BSSs.

The AP virtually connects to each data zone independently. Each data zone's network change or failure does not affect the management of an AP from other data zones. The data zone can configure the AP separately and the AP will apply each configuration. However, if the primary zone goes down, then all the data zones will be affected including the traffic on the data zone.

For example, the first zone has SSID-1, SSID-2 configured and has stand-alone setup, while the second zone has SSID-3, SSID-4 configured and has cluster setup. Then, the MultiZone AP receives both configurations and provides service for all the four SSIDs with no communication between the managed devices.

The MultiZone feature allows the client traffic of different ESS to go to different managed devices into various zones without cross-contamination. The client traffic of the specific ESS is encrypted and tunneled directly from AP to the managed devices using the tunnel mode. All devices in the path including the primary managed device managing the AP are automatically secured. Client wireless frames are encrypted or decrypted for the corresponding SSID data zone managed device in the secure zone.

All the zones can have a maximum of 12 managed devices and 16 VAPs per radio and a maximum of 5 zones are supported including the primary zone.

Starting from ArubaOS 8.3.0.0, MultiZone supports Decrypt Tunnel forwarding mode on the data zone Virtual APs.

Following sections describe the functional flow, licenses, and features of MultiZone:

Functional Flow of a MultiZone AP

The functional flow of a MultiZone AP is as follows:

- AP boots up and terminates on primary zone.
- Receives configuration from primary zone and apply.
- Simultaneously, it connects to each IP address of data zone configured in the MultiZone profile.
- Receives VAP configuration from data zone and apply.
- If common configuration like radio or channel is changed on primary zone, data zone needs to rebootstrap to update.
- If the CPsec is enabled, each data zone managed device should have the AP appropriately white-listed.

Important Points

- CPsec is not mandatory for MultiZone.
- If High Availability is enabled, MultiZone cannot be configured.
- Tunnel mode and Decrypt Tunnel mode are the supported Forward-Modes.
- The primary zone and data zone managed devices do not require to be on the same layer 2 subnet, but, should be layer 3 reachable.
- For the WebCC feature to work seamlessly, the feature should be enabled in each data zone managed device.



The data zone AP ignores the configuration that can affect other zone's BSSs like radio configurations.

Licenses for MultiZone

Starting from ArubaOS 8.2.0.0, data zone managed device will not consume any license and only the primary zone managed device will consume licenses, including the WebCC licenses. Prior to ArubaOS 8.2.0.0, APs connected to data zone managed device consumed PEFNG license, although the data zone managed device still requires PEFNG licenses.

Also, once the AP comes up, the managed device checks if the RFP license was acquired by the AP on the primary zone and Datazone managed device. If not, MultiZone will be disabled on that AP.



The `show ap license-usage` will not count licenses on the data zone managed device for APs that connect to it as a data zone AP.

Hybrid CPsec, Mesh AP, and Mobility Controller Virtual Appliance Support

Starting from ArubaOS 8.2.0.0, hybrid CPsec is supported. That is, CPsec can be enabled or disabled independently for each zone.

Starting from ArubaOS 8.2.0.0, MultiZone is supported for Mobility Controller Virtual Appliance with CPsec enabled. Therefore, a combination of hardware controllers and Mobility Controller Virtual Appliance are supported.

Starting from ArubaOS 8.2.0.0, Mesh is supported on MultiZone only for IPv4.

AP LACP Support for MultiZone

Striping LMS IP can no longer be used to stripe the traffic as the AP has GRE tunnels to more than one managed device. Therefore, starting from ArubaOS 8.2.0.0, LACP is used to stripe traffic on a per UAC basis. That is, the clients or users on the same AP are steered to different UACs and traffic is striped to the UACs.

When MultiZone is enabled, the Striping LMS IP will not be sent to AP. The striping of traffic for the Ethernet interfaces is according to the UAC node.

Limitations

Primary zone managed device is not using striping LMS when the data zone managed device is down.

Client Match Support for MultiZone

Starting from ArubaOS 8.3.0.0, the ClientMatch features like sticky-client and band steering is supported in a MultiZone deployment for Campus APs. ClientMatch in each zone functions independently by controlling clients that are associated to the Virtual APs owned by that zone.

Key Considerations

- ESSIDs must be unique across zones. The same VAP ESSID should not be configured in more than one zone, as this can cause issues in client steering if the zones are co-located.
- ClientMatch configuration in ARM profiles should be set to the default values to ensure that the ClientMatch configuration is common across primary zone and data zones.
- ClientMatch is enabled by default, which is the recommended setting, unless all APs in the MultiZone deployment have the same primary zone.
- ClientMatch spectrum load balancing does not work on radios that are hosting Virtual APs from more than one zone.

RSDB and Dual 5G Bands Support for MultiZone

Starting from ArubaOS 8.3.0.0, MultiZone supports RSDB (Real simultaneous dual band) on AP-203R, AP-203RP and AP-203H access points. Also, MultiZone supports Dual 5 GHz on AP-344 and AP-345 access points.

This feature helps the Data zone managed device to get the current RSDB and dual 5 GHz mode information of the AP and adjust the radio and Virtual AP configurations based on the information.

This feature also ensures that the AP-203R, AP-203RP and AP-203H work in a MultiZone deployment with different RSDB modes and the AP-344 and AP-345 work in MultiZone with different Dual 5 GHz modes.

On a Data zone managed device, execute the following command to display the RSDB and Dual 5 GHz mode:

```
(host) [mynode] #show ap active
```

```
Active AP Table
---------------
Name             Group      IP Address    AP Type  Flags   Uptime Outer IP  Radio 0
Band Ch/EIRP/MaxEIRP/Clients  Radio 1 Band Ch/EIRP/MaxEIRP/Clients
----             -----      ----------    -------  -----   ------ --------  ----------
-------------------------  ------------------------------------
```

```
AP203H-veriwave  rsdb       10.16.140.196  203H     A2aVf  9m:9s   N/A      AP:2.4GHz-
HT:11/6.5/23.5/0
AP203H-veriwave  rsdb       10.16.140.196  203H     A2aUf  9m:34s  N/A      AP:5GHz-
VHT:161E/6.5/23.5/0
AP203H-veriwave  rsdb       10.16.140.196  203H     A2aTf  11m:7s  N/A
snipAP:5GHz-VHT:161E/17.0/19.5/0      AP:2.4GHz-HT:11/12.0/20.5/0

 U = Flex Radio Mode is 5GHz; V = Flex Radio Mode is 2.4GHz; T = Flex Radio Mode is
2.4GHz+5GHz
```

Configuring MultiZone

The primary zone can configure MultiZone profiles to enable the MultiZone feature. The data zone APs are referred to as zone APs. In the data zone, the APs cannot be rebooted, provisioned, or upgraded.

Starting from ArubaOS 8.4.0.0, you can configure either or both IPv4 and IPv6 addresses in one data zone of an AP MultiZone profile. The AP selects either IPv4 or IPv6 address from the data zone configuration.



The **AP-SYSTEM** profile configured in the data zone is ignored except for the LMS redirect option.

The following procedure describes how to configure MultiZone:

To create a MultiZone:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > AP Groups > AP Groups**.
2. Click on the **AP group name** for which you want to create a MultiZone.
3. In the **AP Groups > <AP Group Name>** tab, Click **MultiZone**.
4. Click **Enable multizone** toggle switch to enable MultiZone in an AP.
5. To add a new MultiZone profile, click + in the **MultiZone Profiles** field.
6. Enter the MultiZone name and click **OK**.
7. To add the primary zone and data zone details, click on the MultiZone profile. The parameters are listed in the following table:

Table 74: *MultiZone profile Parameters*

Parameter	Description
IPV4ADDRESS	Specify the IPv4 address of the AP.
IPV6ADDRESS	Specify the IPv6 address of the AP.
No. of WLANs	Specify the number of WLAN SSIDs. The maximum number of WLANs that can be configured are 16.
No. of CONTROLLERS	Specify the number of managed devices. The maximum number of controllers that can be configured are 12.

8. For primary zone, enter the **No. of WLANs** and **No. of Controllers** in their respective field.
9. For data zone, click + and enter the **IPv4 address**, **IPv6 address**, **No. of WLANs**, and **No. of Controllers**.

10. Click **Submit**.
11. Click **Pending Changes**.
12. In the **Pending Changes** window, select the check box and click **Deploy changes**.

To create a MultiZone profile :

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles > AP**.
2. Click **AP MultiZone**. **AP MultiZone profile: New Profile** is displayed.
3. Click **+** in **AP MultiZone profile** to add a new profile.
4. Enter the name of the profile in the **Profile Name** field.



The data zone managed device configuration should only include the VAP profile and AAA profile for the MultiZone AP Group.

5. Click **+** in the **Data zone controller IP** table.
   a. Enter the name of the zone in the **Zone** field.
   b. Enter the IP address in the **IP** field.
   c. Enter the IPv6 address in the **IPv6** field.
   d. Enter the number of virtual APs in the **Num_vaps** field.
   e. Enter the number of zones in the **Num_nodes** field.
6. Click the **Enable/disable MultiZone** check box to enable or disable the MultiZone profile.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

To attach or detach the profile to an ap-group:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > AP Groups**.
2. Add the MultiZone profile to an AP Group.

The following commands create a MultiZone profile, and set the data zone index, primary zone and controller-ip:

```
(host) [mynode] (config)#ap multizone-profile newMZProfile
(host) [mynode](AP multizone profile "newMZProfile") #datazone 1 controller-ip
10.15.146.3 max-vaps 3 max-nodes 4
(host) [mynode](AP multizone profile "newMZProfile") #datazone 2 controller-ip
10.15.144.3 max-vaps 3 max-nodes 2
(host) [mynode](AP multizone profile "newMZProfile") #datazone 3 controller-ip
10.15.144.8 controller-ipv6 2001:1001::201 max-vaps 3 max-nodes 2
(host) [mynode](AP multizone profile "newMZProfile") #primaryzone max-vaps 3 max-nodes
1
(host) [mynode](AP multizone profile "newMZProfile") #multizone-enable
(host) [mynode](AP multizone profile "newMZProfile") #write memory
```



If a zone is in a cluster configuration, the primary zone configuration must have max-nodes and also, ensure that the number of max-nodes in the zone should be more than or equal to the cluster node.

The following commands attach the profile to ap-group or ap-name:

```
(host) [mynode] (config) #ap-group default
```

```
(host) [mynode] (AP group "default") #ap-multizone-profile newMZProfile
(host) [mynode] (AP group "default") #write memory
```

The following command to displays the MultiZone profile , number of data zones and number of virtual APs available:

```
(host) [mynode] (config) #show ap multizone-profile test
```

```
Multizone Enabled

Multizone Table
--------------
Zone  IP Address   IPv6 Address     Max Vaps Allowed  Max Nodes Allowed   Description
----  ----------   ------------     ----------------- -----------------   -----------
0     N/A          N/A               2                 1                  N/A
2     10.15.144.5  2001:1001::201    3                 1

Number of datazones:1
```

The following CLI command displays the MultiZone configured for a particular AP:

```
(host) [mynode] (config) #show ap debug multizone ap-name testAP
```

```
Multizone Table
---------------
Zone  Configured IP      Serving IP       Max Vaps Allowed  Nodes  Flags
----  -------------      ----------       ----------------  -----  -----
0     10.16.84.10        10.16.84.10        13 (0~12)           1     2
1     2008::abc:90:90::4 2008::abc:90:90::4  3 (0-2)            1     V
Flags: C = Cluster; L = Limited nodes; N = Nodes in other zones; 2 = Using IKE version
2; M = Image mismatch; V = IP version mismatch
Number of datazones:1
```

Chapter 21
Virtual Private Networks

Wireless networks can use VPN connections to further secure wireless data from attackers. Mobility Master can be used as a VPN concentrator that terminates all VPN connections from both wired and wireless clients.

This chapter describes the following topics:

- Planning a VPN Configuration on page 459
- Working with VPN Authentication Profiles on page 463
- Configuring a Basic VPN for L2TP/IPsec on page 465
- Configuring a VPN for L2TP/IPsec with IKEv2 on page 471
- Configuring a VPN for Smart Card Clients on page 476
- Configuring a VPN for Clients with User Passwords on page 477
- Configuring Remote Access VPNs for XAuth on page 479
- Working with Remote Access VPNs for PPTP on page 480
- Working with Site-to-Site VPNs on page 480
- Working with VPN Dialers on page 490

Planning a VPN Configuration

The following VPN types can be configured on the Mobility Master or managed devices:

- **Remote access VPNs:** Remote access VPNs allow hosts, such as telecommuters or traveling employees, to connect to private networks (for example, a corporate network) over the Internet. Each host must run VPN client software, which encapsulates and encrypts traffic, then sends it to a VPN gateway at the destination network. The following remote access VPN protocols are supported by Mobility Master:
  - Layer-2 Tunneling Protocol over IPsec (L2TP/IPsec)
  - PPTP
  - XAUTH IKE/IPsec
  - IKEv2 with Certificates
  - IKEv2 with EAP
- **Site-to-site VPNs:** Site-to-site VPNs allow networks to connect to other networks, such as a corporate network. Unlike a remote access VPN, hosts in a site-to-site VPN do not run VPN client software. All traffic for the other network is sent and received through a VPN gateway, which encapsulates and encrypts the traffic.

Before enabling VPN authentication, you must configure the following:

- The default user role for authenticated VPN clients. See Roles and Policies on page 491 for information about configuring user roles.
- The authentication server group used by the managed device to validate clients. See Authentication

Servers on page 192 for configuration details.

NOTE

A server-derived role, if present, takes precedence over the default user role.

Specify the default user role and authentication server group in the VPN authentication **default** profile, as described in the following sections.

NOTE

ESP Tunnel Mode is the only supported IPsec mode of operation. Mobility Master does not support AH and Transport modes.

The following sections describe the considerations for planning a VPN configuration:

Selecting an IKE protocol

Managed devices running ArubaOS 8.0.0.0 support both IKEv1 and IKEv2 protocols to establish IPsec tunnels. Though both IKEv1 and IKEv2 support the same suite-B cryptographic algorithms, IKEv2 is a simpler, faster, and more reliable protocol than IKEv1.

If your IKE policy uses IKEv2, you should be aware of the following caveats when configuring your VPN:

- Separate PSKs cannot be used for each direction of an exchange; both peers must use the same PSK.
- Mixed authentication between both pre-shared keys and certificates is not supported; each authentication exchange requires a single authentication type. For example, if a client authenticates with a PSK, the managed device must also authenticate with a PSK.
- IKEv2 Authentication Headers and IP Payload Compression Protocol are not supported.
- Non-Aruba devices can fragment the large IKE_AUTH packets using the standards described in the **RFC 7383– IKEv2 message fragmentation** when the Aruba device acts as a responder and not as an initiator.

Understanding Suite-B Encryption Licensing

Aruba managed devices support Suite-B cryptographic algorithms when the Advanced Cryptography license is installed. Table 75 describes the Suite-B algorithms supported by Mobility Master IKE Policies and IPsec tunnels. For further details on configuring a VPN to use Suite-B algorithms, see Configuring a VPN for L2TP/IPsec with IKEv2 on page 471.

Table 75: *Suite-B Algorithms Supported by the ACR License*

IKE Policies	Suite-B for IPsec tunnels
hash: SHA-256-128, SHA-384-192	Encryption: AES-128-GCM, AES-256-GCM
Diffie-Hellman Groups: ECP-256, ECP-384	PFS: ECP-256, ECP-384
Pseudo-Random Function (PRF): HMAC_SHA_256, HMAC_SHA_384	—
Suite-B certificates: ECDSA-256, ECDSA-384	—

NOTE

The ArubaOS hardware supports IKE Suite-B AES-128-GCM and AES-256-GCM encryption. The ArubaOS software performs the IKE Suite-B Diffie-Hellman and Certificate-based signature operations, and hash, PFS, and PRF algorithm functions.

The following VPN clients support Suite-B algorithms when establishing an L2TP/IPsec VPN:

Table 76: *Client Support for Suite-B*

Client Operating System	Supported Suite-B IKE Authentication	Supported Suite-B IPsec Encryption
■ Windows client **NOTE:** Windows client operating system includes Windows XP and later versions.	■ IKEv1 Clients using ECDSA Certificates ■ IKEv1/IKEv2 Clients using ECDSA Certificates with L2TP, PPP, EAP-TLS certificate user-authentication	■ AES-128-GCM ■ AES-256-GCM

The Suite-B algorithms described in Table 75 are also supported by Site-to-Site VPNs between Aruba managed devices, or between an Aruba managed device and a server running Windows 2008 or StrongSwan 4.3.

Working with IKEv2 Clients

Not all clients support both the IKEv1 and IKEv2 protocols. Only the clients in Table 77 support IKEv2 with the following authentication types:

Table 77: *VPN Clients Supporting IKEv2*

Windows Client	StrongSwan 4.3 Client	VIA Client
■ Machine authentication with Certificates ■ User name password authentication using EAP-MSCHAPv2 or PEAP-MSCHAPv2 ■ User smart-card authentication with EAP-TLS / IKEv2 **NOTE:** Windows clients using IKEv2 do not support PSK authentication. **NOTE:** Windows client operating system includes Windows 7 and later versions.	■ Machine authentication with Certificates ■ User name password authentication using EAP-MSCHAPv2 ■ Suite-B cryptographic algorithms	■ Machine authentication with Certificates ■ User name password authentication using EAP-MSCHAPv2 or EAP-GTC ■ EAP-TLS using Microsoft cert repository **NOTE:** VIA clients using IKEv2 do not support PSK authentication.

Support for VIA-Published Subnets

Starting from ArubaOS 8.0.1.0, a new feature is introduced in Mobility Master to support IKEv2 configuration (CFG_SET) payload for VIA clients. This is in conformation with section 3.15 of RFC 5996 applicable for route-based VPNs. This feature is disabled by default.

When this feature is enabled, managed devices can accept CFG_SET message with the INTERNAL_IP4_ SUBNET attribute type. When a managed device receives this message, which consists of an IP address and netmask, it adds an entry to the datapath route table that points to the VIA's inner IP address as the next-hop. The datapath route-cache for the VIA's inner IP will point to the tunnel endpoint associated with the VIA.

Enabling Support for VIA-Published Subnets

You can enable the support for VIA-published subnets using the CLI. The following CLI command enables this feature on the Mobility Master:

```
(host)[mynode] (config) #crypto-local isakmp allow-via-subnet-routes
```

The following CLI command disables this feature on the Mobility Master:

```
(host)[mynode] (config)#no crypto-local isakmp allow-via-subnet-routes
```

Verifying Support for VIA-Published Subnets

The following CLI command verifies if the Mobility Master is configured to accept subnet routes from VIA clients:

```
(host)[mynode] #show crypto-local isakmp allow-via-subnet-routes
Controller will accept subnet routes from via client
```

Limitations

The following limitations are applicable to the CFG_SET support feature for Mobility Master:

- This feature supports only IPv4
- This feature is only applicable with IKEv2

For details about how to configure and run VIA on Linux platform, refer to the *VIA 2.3.1 Linux Edition Release Notes*.

Understanding Supported VPN AAA Deployments

If you want to simultaneously deploy various combinations of a VPN client, RAP-psk, RAP-certs, and Campus APs on the same managed device, see Table 78. Each row in this table specifies the allowed combinations of AAA servers for simultaneous deployment. Configuration rules include the following:

- RAP-certs can only use LocalDB-AP.
- An RAP-psk and RAP-cert can only terminate on the same managed device if the VPN profile of the Remote AP uses local database.
- If an RAP-psk is using an external AAA server, the RAP-cert cannot be terminated on the same managed device.
- Clients can use any type of AAA server, regardless of the Remote AP/Campus AP authentication configuration server.

Table 78: *Supported VPN AAA Deployments*

VPN Client	RAP psk	RAP certs	Campus AP
External AAA server 1	LocalDB	LocalDB-AP	CPsec-whitelist
External AAA server 1	External AAA server 1	Not supported	CPsec-whitelist
External AAA server 1	External AAA server 2	Not supported	CPsec-whitelist
LocalDB	LocalDB	LocalDB-AP	CPsec-whitelist
LocalDB	External AAA server 1	Not supported	CPsec-whitelist

Working with Certificate Groups

The certificate group feature allows you to access multiple types of certificates on the same managed device. The following CLI command creates a certificate group:

```
(host) [node] (config) #crypto-local isakmp certificate-group server-certificate
<server_cert-name> ca-certificate <ca_cert-name>
```

The following CLI command displays the certificate groups:

```
(host) [node] #show crypto-local isakmp certificate-group
```

Working with VPN Authentication Profiles

VPN authentication profiles identify an authentication server, the server group to which the authentication server belongs to, and a user-role for authenticated VPN clients. There are three predefined VPN authentication profiles: **default**, **default-rap**, and **default-cap**. These different profiles allow you to use different authentication servers, user roles, and IP pools for VPN, remote AP, and campus AP clients.



You can configure the **default** and **default-rap** profiles, but not the **default-cap** profile.

Table 79: *Predefined Authentication Profile settings*

Parameter	Description	default	default-rap	default-cap
Default Role for authenticated users	The role that is assigned to the authenticated users.	default-vpn-role	default-vpn-role	sys-ap-role 0
Maximum allowed authentication failures	The number of contiguous authentication failures before the station is blacklisted.	0 (feature is disabled)	0 (feature is disabled)	0 (feature is disabled)
Check certificate common name against AAA server	When enabled, this feature verifies that the certificate's common name exists in the server.	disabled	enabled	enabled
Export VPN IP address as a route	When enabled, this feature causes any VPN client address to be exported to OSPF using IPC. **NOTE:** The **Framed-IP-Address** attribute is assigned the IP address as long as the any server returns the attribute. The **Framed-IP-Address** value always has a higher priority than the local address pool.	enabled	enabled	enabled

Parameter	Description	default	default-rap	default-cap
User idle timeout	The user idle timeout value for this profile. Specify the idle timeout value for the client in seconds. Valid range is 30-15300 in multiples of 30 seconds. Enabling this option overrides the global settings configured in the AAA timers. If this is disabled, the global settings are used.	disabled	N/A	N/A
PAN firewalls Integration	Requires IP mapping at Palo Alto Networks firewalls.	disabled	disabled	disabled

The following procedure describes how to modify the **default** VPN authentication profile:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. In the **All Profiles** list, expand **Wireless LAN > VPN Authentication > default** VPN authentication profile.
3. From the **Default Role** drop-down list, select the default user role for authenticated VPN users. (For detailed information on creating and managing user roles and policies, see Roles and Policies on page 491.)
4. (Optional) Set **Max Authentication failures** to an integer value. The default value is 0, which disables this feature.
5. (Optional) If you use client certificates for user authentication, select the **Check certificate common name against AAA server** check box to verify that the certificate's common name exists in the server. This parameter is enabled by default in the **default-cap** and **default-rap** VPN profiles, and is disabled by default on all other VPN profiles.
6. (Optional) Regardless of how an authentication server is contacted, the **Export VPN IP address as a route** option causes any VPN client address to be exported to OSPF using IPC.
7. Enter a **User idle timeout** value, in seconds.
8. (Optional) Enabling **PAN Firewall Integration** requires IP mapping at Palo Alto Networks firewalls. (For more information about PAN firewall integration, see Palo Alto Networks Firewall Integration on page 837.)
9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.
12. In the **All Profiles** list, select the **Server Group** entry below the **Wireless LAN > VPN Authentication > Default** profile.
13. From the **Server Group** drop-down list, select the server group to be used for VPN authentication.
14. Click **Submit**.
15. Click **Pending Changes**.
16. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure VPN authentication:

```
(host) [mm] (config) #aaa authentication vpn default
(host) ^[mm] (VPN Authentication Profile "default") #cert-cn-lookup
(host) ^[mm] (VPN Authentication Profile "default") #clone <source>
(host) ^[mm] (VPN Authentication Profile "default") #default-role <role>
(host) ^[mm] (VPN Authentication Profile "default") #export-route
(host) ^[mm] (VPN Authentication Profile "default") #max-authentication-failures
<number>
(host) ^[mm] (VPN Authentication Profile "default") #pan-integration
(host) ^[mm] (VPN Authentication Profile "default") #radius-accounting <server_group_
name>
(host) ^[mm] (VPN Authentication Profile "default") #server-group <group>
(host) ^[mm] (VPN Authentication Profile "default") #user-idle-timeout <seconds>
```

Configuring a Basic VPN for L2TP/IPsec

The combination of Layer-2 Tunneling Protocol and Internet Protocol Security (L2TP/IPsec) creates a highly-secure technology that enables VPN connections across public networks such as the Internet. L2TP/IPsec provides a logical transport mechanism on which to transmit PPP frames, tunneling, or encapsulation, so that the PPP frames can be sent across an IP network. L2TP/IPsec relies on the PPP connection process to perform user authentication and protocol configuration. With L2TP/IPsec, the user authentication process is encrypted using the Data Encryption Standard (DES) or Triple DES (3DES) algorithm.

L2TP/IPsec using IKEv1 requires two levels of authentication:

- Computer-level authentication with a pre-shared key to create the IPsec SAs to protect the L2TP-encapsulated data.
- User-level authentication through a PPP-based authentication protocol using passwords, SecureID, digital certificates, or smart cards after successful creation of the SAs.



Note that only Windows 7 (and later versions), StrongSwan 4.3, and VIA clients support IKEv2. For additional information on the authentication types supported by these clients, see Working with IKEv2 Clients on page 461.

The following procedure describes how to configure a remote access VPN for L2TP IPsec for clients using pre-shared keys, certificates, or EAP for authentication.

- Defining Authentication Method and Server Addresses on page 465
- Defining Address Pools on page 466
- Enabling Source NAT on page 466
- Selecting Certificates on page 467
- Defining IKEv1 Shared Keys on page 467
- Configuring IKE Policies on page 468
- Setting the IPsec Dynamic Map on page 469

Defining Authentication Method and Server Addresses

The following procedure describes how to define the authentication method and server addresses on Mobility Master:

1. Define the authentication method and server addresses.
2. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab.

3. Expand **IKEv1**.
4. To enable L2TP, select the **L2TP** check box.
5. Select an authentication method for IKEv1 clients. Currently, supported methods include:
   - Password Authentication Protocol (PAP)
   - Extensible Authentication Protocol (EAP)
   - Challenge Handshake Authentication Protocol (CHAP)
   - Microsoft Challenge Handshake Authentication Protocol (MSCHAP)
   - Microsoft Challenge Handshake Authentication Protocol version 2 (MSCHAPv2)
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.
9. Expand **General VPN**. Configure the IP addresses of the **Primary DNS server**, **Secondary DNS server**, **Primary WINS server**, and **Secondary WINS Server** that are pushed to the VPN client.
10. Click **Submit**.
11. Click **Pending Changes**.
12. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Defining Address Pools

The following procedure describes how to define the pool from which the clients are assigned addresses:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand **General VPN**.
3. In the **Address Pools** table, click **+** to open the **Add New Address Pool** section.
4. Specify the **Pool name**, **Start address IPv4 or v6**, and **End address IPv4 or v6**.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

RADIUS Framed-IP-Address for VPN Clients

IP addresses are usually assigned to VPN clients from configured local address pools. However, the **Framed-IP-Address** attribute that is returned from a RADIUS server can be used to assign the address.

VPN clients use different mechanisms to establish VPN connections with Mobility Master, such as IKEv1, IKEv2, EAP, or a user certificate. Regardless of how the RADIUS server is contacted for authentication, the **Framed-IP-Address** attribute is assigned the IP address as long as the RADIUS server returns the attribute. The **Framed-IP-Address** value always has a higher priority than the local address pool.

Enabling Source NAT

The following procedure describes how to enable source NAT on Mobility Master:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand **General VPN**.
3. Select the **Source-NAT** check box if the IP addresses of clients must be translated to access the network.

4. (Optional) If you enable source NAT, select an existing NAT pool from the **NAT pool** drop-down list.

Selecting Certificates

If you are configuring a VPN to support machine authentication using certificates, define the IKE Server certificates for VPN clients using IKE. Note that these certificates must be imported into Mobility Master, as described in Management Access on page 938. The following procedure describes how to select certificates:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand **General VPN**.
3. From the **Server-certificate for VPN clients** drop-down list, select the server certificate for client machines.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.
7. If you are configuring a VPN to support clients using certificates, you must also assign one or more trusted CA certificates to VPN clients.
    a. Expand **Certificates for VPN Clients**.
    b. In the **CA Certificate Assigned for VPN-Clients** table, click **+** to open the **Add New Certificate** section.
    c. Select a **CA certificate** from the drop-down list.
    d. Click **Submit**.
    e. In the **Certificate Groups for VPN-Clients** table, click **+** to open the **Add New Certificate** section.
    f. Select a **Server certificate** and **CA certificate** from the respective drop-down list.
    g. Click **Submit**.
    h. Repeat steps **b** through **g** to add more certificates.
    i. Click **Pending Changes**.
    j. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Defining IKEv1 Shared Keys

If you are configuring a VPN to support IKEv1 and clients using pre-shared keys, you can configure a global IKE key or IKE key for each subnet. Make sure that this key matches the key on the client. The following procedure describes how to define IKEv1 shared keys:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand **Shared Secrets**.
3. In the **IKE Shared Secrets** table, click **+** to open the **Create IKE Group** section.
4. Enter the **Subnet** and **Subnet mask**. To make the IKE key global, enter 0.0.0.0 for both values.
5. Select the **Representation type** from the drop-down list.
6. Enter **Shared key** and repeat it in the **Retype shared key** field.
7. Click **Submit**.

8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Configuring IKE Policies

ArubaOS contains several predefined default IKE policies, as described in the *Default IKE Policy Settings* table. If you do not want to use any of these predefined policies, you can use the procedure below to delete a factory-default policy, edit an existing policy, or create your own custom IKE policy instead.



The IKE policy selections, along with any preshared key, must 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 Aruba dialer is used, these configurations must be made on the dialer prior to downloading the dialer onto the local client.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand **IKEv1**.
3. In the **IKEv1 Policies** table, click an existing policy to edit it, or click **+** to create a new policy.
4. In **Priority**, enter a priority number for this policy. Enter 1 for the configuration to take priority over the default setting.
5. Select the **Enable Policy** check box to enable the policy when it is saved.
6. From the **Encryption** drop-down list, select one of the following encryption types:
   - DES
   - 3DES
   - AES128
   - AES192
   - AES256
7. From the **Hash algorithm** drop-down list, select one of the following hash types:
   - md5
   - sha
   - sha1-96
   - sha2-256-128
   - sha2-384-192
8. ArubaOS VPNs support client authentication using pre-shared keys, RSA digital certificates, or Elliptic Curve Digital Signature Algorithm (ECDSA) certificates. To set the authentication type for the IKE rule, from the **Authentication** drop-down list, select one of the following options:
   - pre-share (for IKEv1 clients using pre-shared keys)
   - rsa-cig (for clients using certificates)
   - ecdsa-256 (for clients using certificates)
   - ecdsa-384 (for clients using certificates)
9. Diffie-Hellman is a key agreement algorithm that allows two parties to agree upon a shared secret, and is used within IKE to securely establish session keys. To set the Diffie–Hellman Group for the ISAKMP policy, from the **Diffie-Hellman group** drop-down list, select one of the following options:
   - Group 1: 768-bit Diffie–Hellman prime modulus group
   - Group 2: 1024-bit Diffie–Hellman prime modulus group
   - Group 14: 2048-bit Diffie–Hellman prime modulus group

- Group 19: 256-bit random Diffie–Hellman ECP modulus group
- Group 20: 384-bit random Diffie–Hellman ECP modulus group



Configuring Diffie–Hellman Group 1 and Group 2 types are not permitted if FIPS mode is enabled.

10. In **Lifetime**, enter a value in the range of 300-86400 seconds to define the lifetime of the security association. The default value is 7200 seconds.
11. Click **Submit**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Setting the IPsec Dynamic Map

Dynamic maps enable IPsec SA negotiations from dynamically addressed IPsec peers. ArubaOS has a predefined IPsec dynamic map for IKEv1. If you do not want to use this predefined map, you can use the procedure below to edit an existing map or create your own custom IPsec dynamic map instead.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand **IKEv1**.
3. In **IKEv1 IPsec Dynamic Maps**, click an existing dynamic map to edit it or click **+** to create a new map.
4. In **Priority**, enter a priority number for this map. Negotiation requests for security associations try to match the highest-priority map first. If that map does not match, the negotiation request continues down the list to the next-highest priority map until a match is made.
5. In **Name**, enter a name for the dynamic map.
6. Select the **Dynamic map** check box.
7. (Optional) Configure PFS settings for the dynamic peer by assigning a Diffie-Hellman prime modulus group. PFS group provides an additional level of security by ensuring that the IPsec SA key was not derived from any other key, and therefore, cannot be compromised if another key is broken. In the **PFS group** drop-down list, select one of the following groups:
   - Group 1: 768-bit Diffie–Hellman prime modulus group
   - Group 2: 1024-bit Diffie–Hellman prime modulus group
   - Group 14: 2048-bit Diffie–Hellman prime modulus group
   - Group 19: 256-bit random Diffie–Hellman ECP modulus group
   - Group 20: 384-bit random Diffie–Hellman ECP modulus group
8. In **Transforms**, select an existing transform to edit it, or click **+** to open the **New Transform** window.



To view current configuration settings for an IPsec transform-set, access the CLI and issue the command **crypto ipsec transform-set tag <transform-set-name>**.

9. Enter a name for the transform in the **Name** field.
10. From the **Encryption** drop-down list, select one of the following encryption types:
    - esp-null
    - esp-des
    - esp-aes128

- esp-aes192
- esp-aes256

11. From the **Hash** algorithm drop-down list, select one of the following hash types:
    - esp-md5-hmac
    - esp-sha-hmac
    - esp-null-hmac
12. Click **Submit**.
13. In **Lifetime(seconds)**, enter a value in the range of 300-86400 seconds to define the lifetime of the security association for the dynamic peer. The default value is 7200 seconds.
14. In **Lifetime(kilobytes)**, enter a value in kilobytes to define the lifetime of the security association for the dynamic peer.
15. Click **Submit**.
16. Click **Pending Changes**.
17. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure a remote access VPN for L2TP IPsec:

1. Define the authentication method and server addresses:

```
(host) [mynode] (config) #vpdn group l2tp
enable
client configuration {dns|wins} <ipaddr1> [<ipaddr2>]
```

2. Enable authentication methods for IKEv1 clients:

```
(host) [mynode] (config) vpdn group l2tp ppp authentication {cache-
securid|chap|eap|mschap|mschapv2|pap
```

3. Create address pools:

```
(host) [mynode] (config) #ip local pool <pool> <start-ipaddr> <end-ipaddr>
```

4. Configure source NAT:

```
(host) [mynode] (config) #ip access-list session srcnatuser any any src-nat
pool <pool> position 1
```

5. If you are configuring a VPN to support machine authentication using certificates, define server certificates for VPN clients using IKEv1:

```
(host) [mynode] (config) #crypto-local isakmp server-certificate <cert>
```

6. If you are configuring a VPN to support IKEv1 Clients using pre-shared keys, you can configure a global IKE key by entering **0.0.0.0** for both the address and netmask parameters in the command below, or configure an IKE key for an individual subnet by specifying the IP address and netmask for that subnet:

```
(host) [mynode] (config) #crypto isakmp key <key> address <ipaddr|> netmask
<mask>
```

7. Define IKE Policies:

```
(host) [mynode] (config) #crypto isakmp policy <priority>
encryption {3des|aes128|aes192|aes256|des}
version v1|v2
```

```
authentication {pre-share|rsa-sig|ecdsa-256ecdsa-384}
group {1|2|19|20}
hash {md5|sha|sha1-96|sha2-256-128|sha2-384-192}
lifetime <seconds>
```

Configuring a VPN for L2TP/IPsec with IKEv2

Only clients running Windows 7 (and later versions), StrongSwan 4.3, and Aruba VIA support IKEv2. For additional information on the authentication types supported by these clients, see "Working with IKEv2 Clients on page 461."

The following procedure describes how to configure a remote access VPN for IKEv2 clients using certificates:

- Defining Authentication Method and Server Addresses on page 471
- Defining Address Pools on page 471
- Enabling Source NAT on page 472
- Selecting Certificates on page 472
- Configuring IKE Policies on page 473
- Setting the IPsec Dynamic Map on page 474

Defining Authentication Method and Server Addresses

The following procedure describes how to define the authentication method and server addresses on Mobility Master:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand **IKEv2**.
3. In **EAP passthrough**, select the EAP passthrough for IKEv2 clients. The currently supported methods include:
   - EAP-TLS
   - EAP-PEAP
   - EAP-MSCHAPv2
   - EAP-GTC
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.
7. Expand **General VPN**.
8. Configure the IP addresses of the **Primary DNS server**, **Secondary DNS server**, **Primary WINS server**, and **Secondary WINS server** that are pushed to the VPN client.
9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Defining Address Pools

The following procedure describes how to define the pool from which the clients are assigned addresses:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand **General VPN**.
3. In the **Address Pools** table, click **+** to open the **Add New Address Pool** section.
4. Specify the **Pool Name**, **Start address IPv4 or v6**, and **End address IPv4 or v6**.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Enabling Source NAT

The following procedure describes how to enable source NAT on Mobility Master:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand **General VPN**.
3. Select the **Source-NAT** check box if the IP addresses of clients must be translated to access the network.
4. (Optional) If you enable source NAT, select an existing NAT pool from the **NAT pool** drop-down list.

Selecting Certificates

If you are configuring a VPN to support machine authentication using certificates, define the IKE Server certificates for VPN clients using IKEv2. Note that these certificate must be imported into Mobility Master, as described in Management Access on page 938.The following procedure describes how to select certificates:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand **General VPN**.
3. From the **Server-certificate for VPN clients** drop-down list, select the server certificate for client machines.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.
7. If you are configuring a VPN to support clients using certificates, you must also assign one or more trusted CA certificates to VPN clients:
    a. Expand **Certificates for VPN Clients**.
    b. In the **CA Certificate Assigned for VPN-Clients** table, click **+** to open the **Add New Certificate** section.
    c. Select a **CA certificate** from the drop-down list.
    d. Click **Submit**.
    e. In the **Certificate Groups for VPN-Clients** table, click **+** to open the **Add New Certificate** section.
    f. Select a **Server certificate** and **CA certificate** from the respective drop-down list.
    g. Click **Submit**.
    h. Repeat steps **b** through **g** to add more certificates.

i. Click **Pending Changes**.

j. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Configuring IKE Policies

ArubaOS contains several predefined default IKE policies, as described in the *Default IKE Policy Settings* table. If you do not want to use any of these predefined policies, you can use the procedures below to delete a factory-default policy, edit an existing policy, or create your own custom IKE policy instead.



The IKE policy selections must 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 Aruba dialer is used, these configurations must be made on the dialer prior to downloading the dialer onto the local client.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand **IKEv2**.
3. In the **IKEv2 Policies** table, click an existing policy to edit it, or click **+** to create a new policy.
4. In **Priority**, enter a priority number for this policy. Enter 1 for the configuration to take priority over the default setting.
5. Select the **Enable Policy** check box to enable the policy when it is saved.
6. From the **Encryption** drop-down list, select one of the following encryption types:
   - DES
   - 3DES
   - AES128
   - AES192
   - AES256
7. From the **Hash algorithm** drop-down list, select one of the following hash types:
   - md5
   - sha
   - sha2-256-128
   - sha2-384-192
8. ArubaOS VPNs support client authentication using pre-shared keys, RSA digital certificates, or ECDSA certificates. To set the authentication type for the IKE rule, from the **Authentication** drop-down list, select one of the following options:
   - pre-share (for IKEv1 clients using pre-shared keys)
   - rsa-sig (for clients using certificates)
   - ecdsa-256 (for clients using certificates)
   - ecdsa-384 (for clients using certificates)
9. Diffie-Hellman is a key agreement algorithm that allows two parties to agree upon a shared secret, and is used within IKE to securely establish session keys. To set the Diffie–Hellman Group for the ISAKMP policy, from the **Diffie-Hellman group** drop-down list, select one of the following options:
   - Group 1: 768-bit Diffie–Hellman prime modulus group
   - Group 14: 2048-bit Diffie–Hellman prime modulus group

- Group 19: 256-bit random Diffie–Hellman ECP modulus group
- Group 20: 384-bit random Diffie–Hellman ECP modulus group



Configuring Diffie–Hellman Group 1 and Group 2 types are not permitted if FIPS mode is enabled.

10. Set the **PRF** value. This algorithm is an HMAC function used to hash certain values during the key exchange:
    - PRF-HMAC-MD5
    - PRF-HMAC-SHA2
    - PRF-HMAC-SHA256
    - PRF-HMAC-SHA384
11. In **Lifetime**, enter a value in the range of 300-86400 seconds to define the lifetime of the security association. The default value is 28800 seconds.
12. Click **Submit**.
13. Click **Pending Changes**.
14. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Setting the IPsec Dynamic Map

Dynamic maps enable IPsec SA negotiations from dynamically addressed IPsec peers. ArubaOS has predefined IPsec dynamic maps for IKEv2. If you do not want to use these predefined maps, you can use the procedures below to delete a factory-default map, edit an existing map, or create your own custom IPsec dynamic map instead.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Click **IKEv2** to expand that section.
3. In **IKEv1 IPSec Dynamic Maps**, click an existing dynamic map to edit it or click **+** to create a new map.
4. In **Priority**, enter a priority number for this map. Negotiation requests for security associations try to match the highest-priority map first. If that map does not match, the negotiation request continues down the list to the next-highest priority map until a match is made.
5. In **Name**, enter a name for the dynamic map.
6. Select the **Dynamic map** check box.
7. (Optional) Configure Perfect Forward Secrecy (PFS) settings for the dynamic peer by assigning a Diffie-Hellman prime modulus group. PFS group provides an additional level of security by ensuring that the IPsec SA key was not derived from any other key, and therefore, cannot be compromised if another key is broken. In the **PFS group** drop-down list, select one of the following groups:
    - Group 1: 768-bit Diffie–Hellman prime modulus group
    - Group 2: 1024-bit Diffie–Hellman prime modulus group
    - Group 14: 2048-bit Diffie–Hellman prime modulus group
    - Group 19: 256-bit random Diffie–Hellman ECP modulus group
    - Group 20: 384-bit random Diffie–Hellman ECP modulus group
8. In **Transforms**, select an existing transform to edit it, or click **+** to open the **New Transform** section.



To view current configuration settings for an IPsec transform-set, access the CLI and issue the command **crypto ipsec transform-set tag <transform-set-name>**.

9. From the **Encryption** drop-down list, select one of the following encryption types:
   - esp-null
   - esp-des
   - esp-3des
   - esp-aes128
   - esp-aes256
10. From the **Hash** algorithm drop-down list, select one of the following hash types:
   - esp-md5-hmac
   - esp-sha-hmac
   - esp-null-hmac
11. Click **Submit**.
12. In **Lifetime(seconds)**, enter a value in the range of 300-86400 seconds to define the lifetime of the security association for the dynamic peer. The default value is 7200 seconds.
13. In **Lifetime(kilobytes)**, enter a value in kilobytes to define the lifetime of the security association for the dynamic peer.
14. Click **Submit**.
15. Click **Pending Changes**.
16. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure remote access VPN for IKEv2 clients using certificates:

To configure a remote access VPN for L2TP IPsec using IKEv2:

1. Define the server addresses:

```
(host) [mynode] (config) #vpdn group l2tp
enable
client configuration {dns|wins} <ipaddr1> [<ipaddr2>]
```

2. Enable authentication methods for IKEv2 clients:

```
(host) [mynode] (config) #crypto isakmp eap-passthrough {eap-gtc|eap-mschapv2|eap-peap|eap-tls}
```

3. Create address pools:

```
(host) [mynode] (config) #ip local pool <pool> <start-ipaddr> <end-ipaddr>
```

4. Configure source NAT:

```
(host) [mynode] (config) #ip access-list session srcnat user any any src-nat pool <pool> position 1
```

5. If you are configuring a VPN to support machine authentication using certificates, define server certificates for VPN clients using IKEv2:

```
(host) [mynode] (config) #crypto-local isakmp server-certificate <cert>
```

NOTE

The IKE PSK value must be between 6-64 characters. To configure a pre-shared IKE key that contains non-alphanumeric characters, surround the key with quotation marks.
For example: **crypto-local isakmp key "key with spaces" fqdn-any**.

6. Define IKEv2 Policies:

```
(host) [mynode] (config) #crypto isakmp policy <priority>
encryption {3des|aes128|aes192|aes256|des}
version v2
authentication {pre-share|rsa-sig|ecdsa-256ecdsa-384}
group {1|2|19|20}
hash {md5|sha|sha1-96|sha2-256-128|sha2-384-192}
prf PRF-HMAC-MD5|PRF-HMAC-SHA1|PRF-HMAC-SHA256|PRF-HMAC-SHA384
lifetime <seconds>
```

7. Define IPsec Tunnel parameters:

```
(host) [mynode] (config) #crypto ipsec
mtu <max-mtu>
transform-set <transform-set-name> esp-3des|esp-aes128|esp-aes128-gcm|esp-aes192|esp-aes256|esp-aes256-gcm|esp-des esp-md5-hmac|esp-null-mac|esp-sha-hmac
```

Configuring a VPN for Smart Card Clients

This section describes how to configure a remote access VPN on a managed device for Microsoft L2TP/IPsec clients with smart cards, which contain a digital certificate allowing user-level authentication without the user entering a username and password. L2TP/IPsec requires two levels of authentication, IKE SA (machine) authentication and user-level authentication with an IKEv2 or PPP-based authentication protocol.

Microsoft clients running Windows 7 (and later versions) support both IKEv1 and IKEv2. Microsoft clients using IKEv2 support machine authentication using RSA certificates (but not ECDSA certificates or pre-shared keys) and smart card user-level authentication with EAP-TLS over IKEv2.

NOTE

Windows 7 (and later version) clients without smart cards also support user password authentication using EAP-MSCHAPv2 or PEAP-MSCHAPv2.

Working with Smart Card clients using IKEv2

To configure a VPN for Windows 7 (and later version) clients using smart cards and IKEv2, follow the procedure described in Configuring a VPN for L2TP/IPsec with IKEv2 on page 471, and ensure that the following settings are configured:

- **L2TP** is enabled
- User authentication is set to **EAP-TLS**
- The IKE policy is configured for **ECDSA** or **RSA** certificate authentication

Working with Smart Card Clients using IKEv1

Microsoft clients using IKEv1, including clients running Windows Vista or earlier versions of Windows, only support machine authentication using a PSK. In this scenario, user-level authentication is performed by an external RADIUS server using PPP EAP-TLS, and client and server certificates are mutually authenticated during the EAP-TLS exchange. During the authentication, EAP-TLS messages from the client are encapsulated into RADIUS messages and forwarded to the server.

Configure the L2TP/IPsec VPN with EAP as the PPP authentication and IKE policy for preshared key authentication of the SA.



On the RADIUS server, configure a remote access policy to allow EAP authentication for smart card users and select a server certificate. The user entry in Microsoft Active Directory must be configured for smart cards.

To configure an L2TP/IPsec VPN for clients using smart cards and IKEv1, ensure that the following settings are configured:

1. On a RADIUS server, a remote access policy must be configured to allow EAP authentication for smart card users and to select a server certificate. The user entry in Microsoft Active Directory must be configured for smart cards. (For detailed information on creating and managing user roles and policies, see Roles and Policies on page 491.)
2. Ensure that the RADIUS server is part of the server group used for VPN authentication.
3. Configure other VPN settings as described in Configuring a VPN for L2TP/IPsec with IKEv2 on page 471, while selecting the following options:
   - Select the **L2TP** check box.
   - Select the **EAP** check box for the Authentication Protocol.
   - Define an IKE Shared Secret to be used for machine authentication. (To make the IKE key global, specify 0.0.0.0 and 0.0.0.0 for both subnet and subnet mask.)
   - Configure the IKE policy for **pre-share** authentication.

Configuring a VPN for Clients with User Passwords

This section describes how to configure a remote access VPN on the managed device for L2TP/IPsec clients with user passwords. L2TP/IPsec requires two levels of authentication, IKE SA authentication and 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. User-level authentication is performed by the internal database of the managed device.

Configure the following:

- AAA database entries for username and passwords
- VPN authentication profile, which defines the internal server group and the default role assigned to authenticated clients
- L2TP/IPsec VPN with PAP as the PPP authentication (IKEv1 only).
- (For IKEv1 clients) An IKE policy for preshared key authentication of the SA.
- (For IKEv2 clients) A server certificate to authenticate the managed device to clients and a CA certificate to authenticate VPN clients.

The following procedure describes how to configure L2TP/IPsec VPN for username and password clients:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** tab.
   a. Select **Internal** from the **Server Groups** table, and then select **Internal** from the **Server Group > Internal** table to display entries for the internal database.
   b. Under **Server Group > Internal > Internal > Users** tab, click **+** to add a new user to the internal server group.
   c. Enter the **User Name** and **Password** information for the client.
   d. Select the **Enabled** check box to activate this entry on creation.
   e. Click **Submit**.
   f. Click **Pending Changes**.
   g. In the **Pending Changes** window, select the check box and click **Deploy changes**.
2. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > L3 Authentication** tab.
   a. From the L3 Authentication List, select **VPN Authentication > default > Server Group** .
   b. Select the **internal** server group from the drop-down list.
   c. Click **Submit**.
   d. Click **Pending Changes**.
   e. In the **Pending Changes** window, select the check box and click **Deploy changes**.
3. Navigate to the **Configuration > Services > VPN** tab.
   a. Expand **IKEv1**.
   b. Select the **L2TP** check box to enable L2TP.
   c. Select the **PAP** check box for **Auth Protocols**.
   d. Click **Submit**.
   e. Click **Pending Changes**.
   f. In the **Pending Changes** window, select the check box and click **Deploy changes**.
4. Configure other VPN settings as described in Configuring a VPN for L2TP/IPsec with IKEv2 on page 471, while ensuring that the following settings are selected:
   - In the **Configuration > Services > VPN** page, select the **L2TP** check box.
   - In the **Configuration > Services > VPN** page, select the **PAP** check box as the authentication protocol.

The following CLI commands configure a L2TP/IPsec VPN for username and password clients using IKEv1:

```
(host) [mynode] (config) #vpdn group l2tp
enable
ppp authentication pap
client dns 101.1.1.245

(host) [mynode] (config) #ip local pool pw-clients 10.1.1.1 10.1.1.250
(host) [mynode] (config) #crypto isakmp key <key> address 0.0.0.0 netmask 0.0.00
(host) [mynode] (config) #crypto isakmp policy 1 authentication pre-share
```

Next, issue the following command to configure client entries in the internal database:

```
(host) [mynode] #local-userdb add username <name> password <password>
```

Configuring Remote Access VPNs for XAuth

XAuth is an Internet draft that permits user authentication after IKE Phase 1 authentication. This authentication prompts the user for a username and password, in which user credentials are authenticated with an external RADIUS , an LDAP server, or the internal database of the managed device. Alternatively, the user can start client authentication with a smart card, which contains a digital certificate to verify the client credentials. IKE Phase 1 authentication can be done with either an IKE preshared key or digital certificates.

Configuring VPNs for XAuth Clients using Smart Cards

This section describes how to configure a remote access VPN on Mobility Master for Cisco VPN XAuth clients using smart cards. Smart cards contain a digital certificate, allowing user-level authentication without the user entering a username and password. IKE Phase 1 authentication can be done with either an IKE preshared key or digital certificates; for XAuth clients using smart cards, the smart card digital certificates must be used for IKE authentication. The client is authenticated with the internal database.

Configure the following:

1. Add entries for Cisco VPN XAuth clients to the internal database of the managed device, an external RADIUS, or an LDAP server. For details on configuring an authentication server, see Authentication Servers on page 192.

   For each client, create an entry in the internal database with the entire Principal name (SubjectAltname in X.509 certificates) or Common Name as it appears on the certificate.

2. Verify that the server with the client data is part of the server group associated with the VPN authentication profile.
3. In the **IKEv1** section of the **Configuration > Services > VPN** tab, select the **L2TP** check box.
4. In the **IKEv1** section of the **Configuration > Services > VPN**tab, select the **XAuth** check box.
5. The Phase 1 IKE exchange for XAuth clients can be either **Main Mode** or **Aggressive Mode**. Aggressive Mode condenses the IKE SA negotiations into three packets (versus six packets for Main Mode). In the **Aggressive group name** field of the **Configuration > Services > VPN** tab, **General VPN** section, enter the authentication group name for aggressive mode to associate this setting to multiple clients. Make sure that the group name matches the aggressive mode group name configured in the VPN client software.
6. Configure other VPN settings as described in Configuring a VPN for L2TP/IPsec with IKEv2 on page 471, while ensuring that the following settings are selected:
   - In the **IKEv1** section of the **Configuration > Services > VPN** tab, select the **L2TP** check box.
   - In the **IKEv1** section of the **Configuration > Services > VPN** tab, select the **XAuth** check box.
   - Define an IKE policy to use **RSA** or **ECDSA** authentication.

Configuring a VPN for XAuth Clients Using a Username and Password

This section describes how to configure a remote access VPN on Mobility Master for Cisco VPN XAuth clients using passwords. IKE Phase 1 authentication is done with an IKE preshared key; users are then prompted to enter their username and password, which is verified with the internal database.

Configure the following:

1. Add entries for Cisco VPN XAuth clients to the internal database of the managed device. For details on configuring an authentication server, see Authentication Servers on page 192



For each client, you need to create an entry in the internal database with the entire Principal name (SubjectAltname in X.509 certificates) or Common Name as it appears on the certificate.

2. Verify that the server with the client data is part of the server group associated with the VPN authentication profile.
3. Configure other VPN settings as described in Configuring a VPN for L2TP/IPsec with IKEv2 on page 471, while ensuring that the following settings are selected:
   - In the**IKEv1** section of the **Configuration > Services > VPN** tab, select the **L2TP** check box.
   - In the **IKEv1** section of the **Configuration > Services > VPN** tab, select the **XAuth** check box.
   - The IKE policy must have **pre-share** authentication.

Working with Remote Access VPNs for PPTP

PPTP is an alternative to L2TP/IPsec. Like L2TP/IPsec, PPTP provides a logical transport mechanism using tunneling or encapsulation to send PPP frames 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 are encrypted through MPPE, which uses the RSA RC-4 encryption algorithm. PPTP connections require user-level authentication through a PPP-based authentication protocol (MSCHAPv2 is the currently-supported method).

The following CLI commands configure PPTP:

```
(host) [mynode] (config) #vpdn group pptp
enable
client configuration {dns|wins} <ipaddr1> [<ipaddr2>]
ppp authentication {mschapv2}
(host) [mynode] (config) #pptp ip local pool <pool_name> <pool_start_address> [<pool_
end_address>]
```

Working with Site-to-Site VPNs

Site-to-site VPNs allow sites in different locations to securely communicate with each other over a Layer-3 network such as the Internet. You can use managed device instead of VPN concentrators to connect the sites. You can also use a VPN concentrator at one site and a managed device at the other site.

Mobility Master 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.
- The management MAC address of the Mobility Master should be added as the peer MAC address in the managed device to establish the IKE/IPSEC tunnel with the Mobility Master.For more information on configuring the MAC address for MAC-based PSK authentication, see the Configuring MAC Address for PSK Authenticationsection.
- **Suite-B cryptographic algorithms**: Managed devices support Suite-B cryptographic algorithms when the Advanced Cryptography license is installed. For more information, see Understanding Suite-B Encryption Licensing on page 460.

- **Digital certificates:** You can configure an RSA or ECDSA server certificate and a CA certificate for each site-to-site VPN IPsec map configuration. If you use certificate-based authentication, the peer must be identified by its certificate subject name, distinguished name (for deployments using IKEv2), or by the peer's IP address (for IKEv1). For more information about importing server and CA certificates into Mobility Master, see Management Access on page 938.



Certificate-based authentication is only supported for site-to-site VPN between two managed devices with static IP addresses. IKEv1 site-to-site tunnels cannot be created between a Mobility Master and managed device.

Enable IP compression in an IPsec map to reduce the size of data frames transmitted over a site-to-site VPN between 7200 Series or 7000 Series managed devices using IKEv2 authentication. IP compression can reduce the time required to transmit the frame across the network. When this hardware-based compression feature is enabled, the quality of unencrypted traffic (such as Lync or Voice traffic) is not compromised by increased latency or decreased throughput. IP compression is disabled by default.



This feature is only supported in an IPv4 network using IKEv2. This feature cannot be enabled on a 7205managed device or on a site-to-site VPN established using IKEv1.

Configuring MAC Address for PSK Authentication

On Mobility Master, you can configure the MAC address of the managed device to be used for PSK authentication. The following procedure describes how to configure the MAC address of the managed device for PSK authentication:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Controllers**.
2. Click **+** under **Local Controller IPSec Keys** table.
3. Select **Mac-based PSK** from the **Authentication** drop-down list.
4. Enter the **Mac address**.
5. Enter the **IPSec key**.
6. Retype the IPsec key.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures the MAC address of the managed device for PSK authentication:

```
(host) [mynode] (config) #local-peer-mac 00:0c:29:00:00:00 ipsec 123456
```

You can configure the MAC-based PSK authentication on the managed device.

Working with Third-Party Devices

Managed Devices can use IKEv1 or IKEv2 to establish a site-to-site VPN with another managed device or third-party remote client devices. Devices running Microsoft® Windows 2008 can use Suite-B cryptographic algorithms and IKEv1 to support authentication using RSA or ECDSA. StrongSwan® 4.3 devices can use IKEv2 to support authentication using RSA or ECDSA certificates, Suite-B cryptographic algorithms, and pre-shared keys. These two remote clients are tested to work with managed devices using Suite-B cryptographic algorithm.

Working with Site-to-Site VPNs with Dynamic IP Addresses

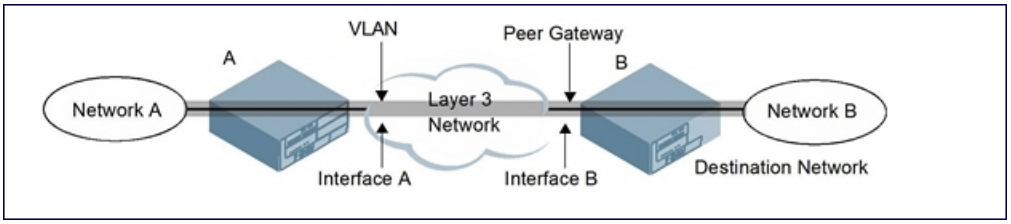
ArubaOS supports site-to-site VPNs with two statically addressed managed devices, or with one static and one dynamically addressed managed device. Two methods are supported to enable dynamically addressed peers:

- **Pre-shared Key Authentication with IKE Aggressive Mode:** The managed device with a dynamic IP address must be configured as the initiator of IKE Aggressive-mode for Site-Site VPNs, while the managed device with a static IP address must be configured as the responder of IKE Aggressive mode. Note that when the managed device is operating in FIPS mode, IKE aggressive mode must be disabled.
- **X.509 certificates:** IPsec peers will identify each other using the subject name of X.509 certificates. IKE operates in main mode when this option is selected. This method is preferred from a security standpoint.

Understanding VPN Topologies

You must configure VPN settings on the managed devices at both the local and remote sites. In the following figure, a VPN tunnel connects Network A to Network B across the Internet.

Figure 83 *Site-to-Site VPN Configuration Components*



To configure the VPN tunnel on managed device A, you must configure the following:

- The source network (Network A)
- The destination network (Network B)
- The VLAN on which managed device A's interface to the Layer-3 network is located (Interface A in Figure 83)
- The peer gateway, which is the IP address of managed device B's interface to the Layer-3 network (Interface B in Figure 83)

NOTE

Configure VPN settings on the managed device at both the local and remote sites.

Configuring Site-to-Site VPNs

The following procedure describes how to configure a site-to-site VPN:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > VPN** tab and expand **Site to Site** .
2. In the **IPsec Maps** section, click **+** to open the **Create New IPsec** section.
3. Enter a name for this VPN connection in the **Name** field.
4. Select the **Enabled** check box so this configuration takes effect as soon as it is saved.
5. In the **Priority** field, enter a priority level for the IPsec map. Negotiation requests for security associations try to match the highest-priority map first. If that map does not match, the

negotiation request continues down the list to the next-highest priority map until a match is made.

6. Select a **Source network type** to specify whether the VPN *source*, the local network connected to the managed device, is defined by an IP address or a VLAN ID.
   - If you selected **IP Address**, enter the IP address and netmask for the source network (see managed device A in Figure 83).
   - If you selected **VLAN**, click the **VLAN** drop-down list and select the VLAN ID for the source network.
7. In the **Destination network** and **Destination subnet mask** fields, enter the IP address and netmask for the *destination*, the remote network to which the local network communicates (see managed device B in Figure 83).
8. The **SA Lifetime** parameter defines the lifetime of the security association in seconds and kilobytes. For seconds, the default value is 7200. To change this value, enter a value between 300 and 86400 seconds. Range: 1000–1000000000 kilobytes.
9. Click the **IKE version** drop-down list and select **v1** to configure the VPN for IKEv1, or **v2** for IKEv2.
10. (Optional) Click the **IKE policy** drop-down list and select a predefined or custom IKE policy to apply to the IPsec map. For more information on default IKE policies, see Table 80.
11. IKEv2 site-to-site VPNs between Mobility Master and 7000 Series managed devices support traffic compression between those devices. Select the **IP compression** check box to enable compression for traffic in the site-to-site tunnel.
12. Select the **Factory certificate authentication** check box to enable the authentication.
13. Select the **VLAN** containing the interface of the managed device that connects to the Layer-3 network (see Interface A in Figure 83). This determines the source IP address used to initiate IKE.
14. If you enable **PFS** mode, new session keys are not derived from previously used session keys. Therefore, if a key is compromised, that compromised key does not affect any previous session keys. PFS mode is disabled by default. To enable this feature, click the **PFS** drop-down list and select one of the following **Perfect Forward Secrecy** modes:
    - **group1:** 768-bit Diffie–Hellman prime modulus group
    - **group2:** 1024-bit Diffie–Hellman prime modulus group
    - **group14:** 2048-bit Diffie–Hellman prime modulus group
    - **group19:** 256-bit random Diffie–Hellman ECP modulus group
    - **group20:** 384-bit random Diffie–Hellman ECP modulus group
15. Select the **Pre-connect** check box to establish the VPN connection, even if there is no traffic being sent from the local network. If you do not select this, the VPN connection is established only when traffic is sent from the local network to the remote network.
16. Select the **Trusted tunnel** check box if traffic between the networks is trusted. If you do not select this, traffic between the networks is untrusted.



Ensure that you always enable the **Trusted tunnel** option. The traffic cannot pass through if this option is disabled.

17. Select the **Enforce NAT-T** check box to enforce UDP 4500 for IKE and IPsec. This option is disabled by default.
18. Add one or more transform sets to be used by the IPsec map. Click +, and select an existing transform set or create a new one. Then click **Submit** to add that transform set to the IPsec map.

19. For site-to-site VPNs with dynamically addressed peers, select **Dynamic** from the **Remote peer addressing** drop-down list.
    a. From the **Peer gateway** drop-down list, select **Initiator** if the dynamically addressed switch is the *initiator* of IKE Aggressive-mode for Site-Site VPNs, or select **Responder** if the dynamically addressed switch is the *responder* for IKE Aggressive-mode.
    b. In the **FQDN** field, enter a FQDN for the managed device. If the managed device is defined as a dynamically addressed responder, you can select **All Peers** to make the managed device a responder for all VPN peers, or select **Per Peer Id** and specify the FQDN to make the managed device a responder for one specific initiator.
20. For **Remote peer addressing** that is **Static**, select one of the supported peer gateway types:
    - **IP Address**: Select this option to identify the remote end point of the VPN tunnel using an IP address.
    - **FQDN** : This option allows you to use same FQDN across different branches. The FQDN resolves to different IP addresses for each branch, based on its local DNS setting.
21. Define the Peer Gateway using an IP address or FQDN.
    - If you use IKEv1 to establish a site-to-site VPN for a statically addressed remote peer and selected **IP Address** in the previous step, enter the IP address of the interface used by the remote peer to connect to the L3 network in the **Peer gateway IPv4 or v6** field (see Interface B in Figure 83).
    - If you are configuring an IPsec map for a dynamically addressed remote peer, and selected **IP Address** in the previous step, leave the **Peer gateway IPv4 or v6** set to its default value of 0.0.0.0.
    - If you selected **FQDN** as the peer gateway type in the previous step, enter the fully qualified domain name for the remote peer.
22. Select one of the following authentication types:
    a. For PSK authentication, select **PSK**, select the **Representation type**, then enter a shared secret in the **IKE shared secret** and **Retype shared secret** fields. This authentication type is generally required in IPsec maps for a VPN with dynamically addressed peers, but can also be used for a static site-to-site VPN.
    b. For certificate authentication, select **Certificate**, then click the **Server certificate** and **CA certificate** drop-down lists to select certificates previously imported into the managed device. See Management Access on page 938 for more information. Enter the **Peer certificate subject name**.



To identify the subject name of a peer certificate, issue the following command in the CLI:

```
Show crypto-local pki servercert <certname> subject
```

23. Click **Submit**.
24. Click **Pending Changes**.
25. In the **Pending Changes** window, select the check box and click **Deploy changes**.
26. Click the **IKEv1** or **IKEv2** section (match the IKE version that you selected in Step 9) to configure an IKE policy.
    a. Under **IKE Policies**, click **+** to open the **Add IKE Policy** configuration page.
    b. Set the **Priority** to **1** for this configuration to take priority over the Default setting.

c. Select the **Enable policy** check box so the configuration takes effect as soon as it is saved.

d. Set the **Encryption** from the drop-down list.

e. Set the **HASH algorithm** from the drop-down list.

f. Set the **Authentication** to **pre-share** if you use pre-shared keys. If you use certificate-based IKE, select **rsa** or **ecdsa**.

g. Set the **Diffie-Hellman group** from the drop-down list.

h. Set the **Lifetime** to define the lifetime of the security association in seconds. The default value is 28800 seconds. To change this value, enter a value between 300 and 86400 seconds.

i. The IKE policy selections, including any PSK, must be reflected in the VPN client configuration. When using a third-party VPN client, set the VPN configuration on clients to match the choices made above. If you use the Aruba dialer, you must configure the dialer prior to downloading the dialer onto the local client.

j. Click **Submit**.

k. Click **Pending Changes**.

l. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure a site-to-site VPN with two static IP managed devices using IKEv1:

```
(host) [mynode] (config) #crypto-local ipsec-map <ipsec-map-name> <ipsec-map-number>
src-net <ipaddr> <mask>
dst-net <ipsec-map-dst-net> <ipsec-map-dst-mask>
peer-ip <ipaddr>
vlan <ipsec-map-vlan-id>
version {v1|v2}
peer-cert-dn <peer-dn>
pre-connect {enable|disable}
trusted enable
```



The *trusted <disable>* sub-parameter is not supported on the managed device. You must always use the *trusted <enable>* sub-parameter so that the traffic can pass through.

For certificate authentication:

```
set ca-certificate <cacert-name>
set server-certificate <cert-name>

(host) [mynode] (config) #crypto isakmp policy <priority>
encryption {3DES|AES128|AES192|AES256|DES}
version {v1|v2}
authentication {pre-share|rsa-sig|ecdsa-256|ecdsa-384}}
group {1|2|14|19|20}
hash {md5|sha|sha1-96|sha2-256-128|sha2-384-192}
lifetime <seconds>
```

For PSK authentication:

```
(host) [mynode] (config) #crypto-local isakmp {key <keystring>|key-hex <keystring>}
address <peer-address> netmask <mask>

(host) [mynode] (config) #crypto isakmp policy <priority>
encryption {3DES|AES128|AES192|AES256|DES}
version {v1|v2}
authentication {pre-share|rsa-sig|ecdsa-256|ecdsa-384}}
group {1|2|14|19|20}
```

```
hash {md5|sha|sha1-96|sha2-256-128|sha2-384-192}
lifetime <seconds>
```

To configure site-to-site VPN with a static and dynamically addressed managed device that initiates IKE Aggressive-mode for Site-Site VPN:

```
(host) [mynode] (config) #crypto-local ipsec-map <ipsec-map-name> <ipsec-map-number>
src-net <ipaddr> <mask>
dst-net <ipsec-map-dst-net> <ipsec-map-dst-mask>
peer-ip <ipaddr>
local-fqdn <local_id_fqdn>
vlan <ipsec-map-vlan-id>
pre-connect {enable|disable}
trusted enable
```

For the Pre-shared-key:

```
(host) [mynode] (config) #crypto-local isakmp {key <keystring>|key-hex <keystring>}
address <peer-address> netmask 255.255.255.255
```

For a static IP managed device that responds to IKE Aggressive-mode for Site-Site VPN:

```
(host) [mynode] (config) #crypto-local ipsec-map <ipsec-map-name 2> <ipsec-map-number>
src-net <ipaddr> <mask>
dst-net <ipsec-map-dst-net> <ipsec-map-dst-mask>
peer-ip 0.0.0.0
peer-fqdn fqdn-id <peer_id_fqdn>
vlan <ipsec-map-vlan-id>
trusted enable
```

For the Pre-shared-key:

```
(host) [mynode] (config) #crypto-local isakmp {key <keystring>|key-hex <keystring>}
fqdn <ike-id-fqdn>
```

For a static IP managed device that responds to IKE Aggressive-mode for Site-Site VPN with one PSK for All FQDNs:

```
(host) [mynode] (config) #crypto-local ipsec-map <ipsec-map-name 2> <ipsec-map-number>

src-net <ipaddr> <mask>
peer-ip 0.0.0.0
peer-fqdn any-fqdn
vlan <ipsec-map-vlan-id>
trusted enable
```

For the Pre-shared-key for All FQDNs:

```
(host) [mynode] (config) #crypto-local isakmp {key <keystring>|key-hex <keystring>}

fqdn-any
```

Supporting Null Encryption for IKEv1

Starting from ArubaOS 8.1.0.0, XLP-based controllers are supported with null encryption for IKEv1 as an encryption algorithm. This helps in reducing the load on the local router for internet destined traffic.

Null encryption does not increase the security of traffic routed but is used only to imply that no encryption method is used over a particular transmission. Null Encryption can now be configured as an encryption algorithm in transform set, which can be used in any crypto map.



Since null encryption is supported only for IKEv1, it should be used only for crypto maps with version 1.

The following procedure describes how to configure a new transformation set with null encryption as the encryption algorithm:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Expand **IKEv1**.
3. In the **IKEv1 IPSec Dynamic Maps** table, click **+** to access the **Add IKEv1 Dynamic Map** section.
4. Click **+** in the **Transforms** field.
5. Select the **Add new transform** option in the **New Transform** window.
6. Select **esp-null** from the **Encryption** drop-down list.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following procedure describes how to add the transformation set in the crypto map created:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Click **Site to Site** accordion.
3. In the **IPSec Maps** table, click **+** to access the **Create New Ipsec** section.
4. Click **+** in the **Transforms** field.
5. Select the **Add existing transform** option in the **New Transform** window.
6. Select an existing transform and click **OK**.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures a new transformation set with null encryption as the encryption algorithm:

```
(host) [mynode] (config) #crypto ipsec transform-set test esp-null esp-sha-hmac
```

The following CLI commands add the transformation set in the crypto map created:

```
(host) [mynode] (config) #crypto-local ipsec-map test_map 500
(host) [mynode] (config-ipsec-map) #set transform-set test
```

Adding ANY-ANY Crypto Map

Starting from ArubaOS 8.1.0.0, any-any selectors are negotiated in IKEv1 to enable the option of having numerous tunnels. After pre-connect flag is enabled for IPsec map, IKE triggers the tunnel to the peer ip and proposes any-any traffic selector.

PBR can also be configured to send specific or all traffic on to the ipsec map and can be applied to any vlan, port, or user role.



Policy Based Routing is required for any-any traffic selector and is supported only for IKEv1.

Data traffic trigger is not supported in ArubaOS 8.1.0.0.

The following procedure describes how to enable crypto map to allow any any traffic selector:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > VPN** tab.
2. Click **Site to Site** accordion.
3. In the **IPSec Maps** table, click **+** to access the **Create New Ipsec** section.
4. Enter a **Name**.
5. Select **Any** from the **Source network type** drop-down list.
6. Select **Any** from the **Destination network type** drop-down list.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands enable crypto map to allow any any traffic selector:

```
(host) [mynode] (config-ipsec-map)#  src-net any
(host) [mynode] (config-ipsec-map)#  dst-net any
```

The following CLI commands configure PBR to send all or specific traffic onto the IPsec map:

```
(host) [mynode] (config) #ip access-list route ipsec-pbr
(host) [mynode] (config-route-ipsec-pbr)#any any any route ipsec-map <ipsec-map-name>
```

The following CLI commands apply PBR to vlan, port, or user role:

```
(host) [mynode] (config) #interface vlan <id>
(host) [mynode] (config-subif) #ip access-group <name> in
```

Dead Peer Detection

DPD is enabled by default on the managed device for site-to-site VPNs. DPD, as described in RFC 3706, “A Traffic-Based Method of Detecting Dead IKE Peers,” uses IPsec traffic patterns to minimize the number of IKE messages required to determine the liveliness of an IKE peer.

After a dead peer is detected, the managed device tears down the IPsec session. Once the network path or other failure condition has been corrected, a new IPsec session is automatically re-established.

The following CLI command configures DPD parameters:

```
(host) [mynode] (config) #crypto-local isakmp dpd idle-timeout <idle_sec> retry-
timeout <retry_sec> retry-attempts <retry_number>
```

About Default IKE Policies

ArubaOS includes the following default IKE policies. These policies are predefined, but can be edited and deleted. You can do this in the CLI by using the **crypto isakmp policy** and **crypto dynamic-map** commands, or the WebUI by navigating to **Configuration > Services > VPN**. To delete an IKE policy, select an existing policy and click the trash icon to delete the policy.

Table 80: *Default IKE Policy Settings*

Policy Name	Policy Number	IKE Version	Encryption Algorithm	Hash Algorithm	Authen-tication Method	PRF Method	Diffie-Hellman Group
Default protection suite	10001	IKEv1	3DES-168	SHA 160	Pre-Shared Key	N/A	2 (1024 bit)
Default RAP Certificate protection suite	10002	IKEv1	AES -256	SHA 160	RSA Signature	N/A	2 (1024 bit)
Default RAP PSK protection suite	10003	IKEv1	AES -256	SHA 160	Pre-Shared Key	N/A	2 (1024 bit)
Default RAP IKEv2 RSA protection suite	1004	IKEv2	AES -256	SHA 160	RSA Signature	hmac-sha1	2 (1024 bit)
Default Cluster PSK protection suite	10005	IKEv1	AES -256	SHA160	Pre-Shared Key	Pre-Shared Key	2 (1024 bit)
Default IKEv2 RSA protection suite	1006	IKEv2	AES - 128	SHA 96	RSA Signature	hmac-sha1	2 (1024 bit)
Default IKEv2 PSK protection suite	10007	IKEv2	AES - 128	SHA 96	Pre-shared key	hmac-sha1	2 (1024 bit)
Default Suite-B 128bit ECDSA protection suite	10008	IKEv2	AES - 128	SHA 256-128	ECDSA-256 Signature	hmac-sha2-256	Random ECP Group (256 bit)
Default Suite-B 256 bit ECDSA protection suite	10009	IKEv2	AES -256	SHA 384-192	ECDSA-384 Signature	hmac-sha2-384	Random ECP Group (384 bit)

Policy Name	Policy Number	IKE Version	Encryption Algorithm	Hash Algorithm	Authen-tication Method	PRF Method	Diffie-Hellman Group
Default Suite-B 128bit IKEv1 ECDSA protection suite	10010	IKEv1	AES-GCM-128	SHA 256-128	ECDSA-256 Signature	hmac-sha2-256	Random ECP Group (256 bit)
Default Suite-B 256-bit IKEv1 ECDSA protection suite	10011	IKEv1	AES-GCM-256	SHA 256-128	ECDSA-256 Signature	hmac-sha2-256	Random ECP Group (256 bit)

Working with VPN Dialers

For Windows clients, a dialer can be downloaded from Mobility Master to auto-configure tunnel settings on the client.

Configuring VPN Dialer

The following CLI commands configure the VPN dialer:

```
(host) [mynode] (config) #vpn-dialer <name>
enable {dnctclear|l2tp|pptp|secureid_newpinmode|wirednowifi}
ike authentication {pre-share <key>|rsa-sig}
ike encryption {3des|des}
ike group {1|2}
ike hash {md5|sha}
ipsec encryption {esp-3des|esp-des}
ipsec hash {esp-md5-hmac|esp-sha-hmac}
ppp authentication {cache-securid|chap|mschap|mschapv2|pap}
```

Assigning a Dialer to a User Role

The VPN dialer can be downloaded using Captive Portal. For the user role assigned through Captive Portal, configure the dialer by using the dialer name.

For example, if the Captive Portal client is assigned to the *guest* role after logging in, and the dialer is called *mydialer*, configure *mydialer* as the dialer to be used in the guest role.

The following CLI command configures the Captive Portal dialer for a user-role:

```
(host) [mynode] (config) #user-role <role>
dialer <name>
```

Roles and Policies

The client in an Aruba user-centric network is associated with a *user role*, which determines the client's network privileges, how often must they be re-authenticated, and which bandwidth contracts are applicable to each client. A *policy* is a set of rules that apply to the traffic that passes through the Aruba managed device. You can specify one or more policies for a user role. Finally, you can assign a user role to clients before or after they get authenticated to the system.

This chapter describes assigning and creating roles and policies for Aruba clients.

The following list displays the key topics discussed:

- Firewall Policies on page 491
- Creating a User Role on page 501
- Workflow for Assigning a User Role on page 504
- Understanding Global Firewall Parameters on page 509
- AppRF 2.0 on page 515



This chapter describes configuring firewall policies and parameters that relate to IPv4 traffic. See IPv6 Support on page 122 for information about configuring IPv6 firewall policies and parameters.

Firewall Policies

A firewall policy identifies specific characteristics about a data packet passing through the Aruba Managed Device and takes some action based on that identification. In an Aruba Managed Device, 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 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 and ACLs have three main functional differences. Firewall policies differ from ACLs in the following ways:

- Firewall policies are *stateful*, meaning that they recognize flows in a network and keep track of the state of sessions. For example, if a firewall policy permits telnet traffic from a client, the policy also recognizes that inbound traffic associated with that session should be allowed.
- Firewall policies are *bi-directional*, meaning that they keep track of data connections traveling into or out of the network. ACLs are normally applied to either traffic inbound to an interface or outbound from an interface.
- Firewall policies are *dynamic*, meaning that address information in the policy rules can change as the policies are applied to users. For example, the alias *user* in a policy automatically applies to the IP address assigned to a particular user. ACLs typically require static IP addresses in the rule.



You can apply IPv4 and IPv6 firewall policies to the same user role. See IPv6 Support on page 122 for information about configuring IPv6 firewall policies.

Workflow for Configuring Firewall Policies

You can configure one or more firewall policies. This section describes how to configure the rules that constitute a firewall policy. In order to configure the correct firewall policies, ensure that you first understand ACL, how to work with ACLs, and what are role-based ACLs.

Working With ACLs

ACLs are a common way of restricting certain types of traffic on a physical port. ArubaOS provides the following types of ACLs:

- Standard ACLs permit or deny traffic based on the source IP address of the packet. Standard ACLS can be either named or numbered, with valid numbers in the range of 1-99 and 1300-1399. Standard ACLs use a bitwise mask to specify the portion of the source IP address to be matched.
- Extended ACLs permit or deny traffic based on source or destination IP address, source or destination port number, or IP protocol. Extended ACLs can be named or numbered, with valid numbers in the range 100-199 and 2000-2699.
- MAC ACLs are used to filter traffic on a specific source MAC address or range of MAC addresses. Optionally, you can mirror packets to a datapath or remote destination for troubleshooting and debugging purposes. MAC ACLs can be either named or numbered, with valid numbers in the range of 700-799 and 1200-1299.
- Ethertype ACLs are used to filter based on the Ethertype field in the frame header. Optionally, you can mirror packets to a datapath or remote destination for troubleshooting and debugging purposes. Ethertype ACLs can be either named or numbered, with valid numbers in the range of 200-299.These ACLs can be used to permit IP while blocking other non-IP protocols, such as IPX or AppleTalk.
- Service ACLs provide a generic way to restrict how protocols and services from specific hosts and subnets to the Mobility Master are used. Rules with this ACL are applied to all traffic on the Mobility Master regardless of the ingress port or VLAN.
- Routing ACLs forward packets to a device defined by an IPsec map, a next-hop list, a tunnel or a tunnel group.



Routing ACL is the only supported ACL type that can be configured on a VLAN Interface. Other ACL types are not supported.

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.

Role-Based ACL

Role-based ACL is a feature available on Aruba controllers to apply policies to traffic matching a particular user role. Earlier this feature was supported only when the users were present in the same controller. Starting from ArubaOS 8.6.0.0, this feature is extended to support multi-controller deployments. Role- to- role ACL can now be assigned to two users terminating on different controllers. This feature can be configured by creating a policy domain group profile and adding the IP address of the controllers.

Role-based ACL supports mix of controller models with the exception of and x86 Virtual Mobility controllers. To apply role-based ACL for 9004 and x86 Virtual Mobility Controllers models, all the

controllers have to be either 9004 or x86 VMCs respectively. To apply role-based ACL to x86 Virtual Mobility Controllers, all the controllers have to be managed by the same Mobility Master.

Role-based ACL works across multiple controllers only if the role is configured as a destination role in at least one ACL.

Role-based ACL cannot be applied to the following:

- L2 multicast traffic
- L3 multicast/broadcast traffic
- ClearPass Policy Manager downloadable user role

The following CLI commands create role-based ACL in a multi-controller deployment:

```
(host) [md] policy-domain group-profile <name>
(host) [md] (Policy Domain Profile "name") controller <ip> <macaddress>
```



Multiple policy domains for group profiles are supported. The command should be executed in the /md node and the policy domain group profile supports IPv4 and IPv6 addresses but a combination of both is not supported.

Limitations

- Each node can be part of one profile only.
- All policy domain profiles can be applied at **/md** nodes only.
- Each policy domain profile can only have either all IPv4 or all IPv6 nodes. Mix of IPv4 and IPv6 nodes are not allowed.
- Managed devices should be part of a single domain. You cannot add a managed device to a Mobility Master, which is already part of another domain.



All managed devices should be running ArubaOS 8.7.0.0 when multiple policy domain manager profiles are configured.

The tasks for configuring a firewall policy are:

1. Configure the rules that constitute in creating a firewall policy.

   See Creating a Firewall Policy on page 493.
2. Create a network alias. A network service alias defines a TCP, UDP, or IP protocol and a list or range of ports supported by that service.

   See Creating a Network Service Alias on page 497
3. Create an ACL whitelist. The ACL white list consists of rules that explicitly permit or deny session traffic from being forwarded to or blocked from the managed device.

   See Creating an ACL White List on page 498
4. Create a local net destination override. This feature provides a scalable solution to create a local net destination override.

   See Override Local Network Destination on page 499

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 81 describes required and optional parameters for a rule.

The following procedure describes how to create a web-only policy that allows web (HTTP and HTTPS) access:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Roles & Policies > Policies** tab.
2. Click **+** to create a new policy.
3. Enter the policy name in the **Policy name** field.
4. Select the policy type from the **Policy type** drop-down list. You can select **Ethertype**, **Extended, MAC, Route, Session,** or **Standard**.
5. Click **Submit**.
6. Select the policy created and click **+** in the **Policy <policy name>** table.
7. Select **Access Control** option in the **Rule Type** field.
8. Click **OK**.
9. To add a rule that allows HTTP traffic.
   a. Under **Service/app**, select **Service** from the drop-down list.
   b. Select **svc-http** from the **Servicealias** drop-down list.
10. Click **Submit**.



Rules can be re-ordered by using the up and down buttons provided for each rule.

11. Click **Submit** to apply this configuration. The policy is not created until the configuration is applied.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command creates a web-only policy that allows web (HTTP and HTTPS) access:

```
(host) [md] (config) #ip access-list session web-only
```

Table 81: *Firewall Policy Rule Parameters*

Parameter	Description
IP version	Specifies whether the policy applies to IPv4 or IPv6 traffic.
Source (required)	Source of the traffic, which can be one of the following: ■ any: Acts as a wildcard and applies to any source address. ■ user: Refers to traffic from the wireless client. ■ host: Refers to traffic from a specific host. When this option is chosen, you must configure the IP address of the host. ■ network: Refers to a traffic that has a source IP from a subnet of IP addresses. When this option is chosen, you must configure the IP address and network mask of the subnet. ■ alias: Refers to using an alias for a host or network. You configure the alias by navigating to the **Configuration > Roles & Policies > Policies** tab. Select a policy created and click **+** to create a Rule. Select the **Access Control** option in the **Rule Type**. Select **Alias** from the **Destination** drop-down list and the alias name from the **Destination alias** drop-down list. Select a Source from the traffic **Source**

Parameter	Description
	drop-down list.
Destination (required)	Destination of the traffic, which can be configured in the same manner as Source.
Service/app (required)	Type of traffic, which can be one of the following: ■ any: This option specifies that this rule applies to any type of traffic. ■ application: For session and route policies on a 7000 Series managed device, you can create a rule that applies to a specific application type. Click the **Application** drop-down list and select an application type. ■ web category/ reputation: For session policies on a 7000 Series managed device, you can create a rule that applies to a specific web category or application type. For more information on web category classification, see Traffic Analysis on page 430 ■ tcp: Using this option, you configure a range of TCP port(s) to match for the rule to be applied. ■ udp: Using this option, you configure a range of UDP port(s) to match for the rule to be applied. ■ service: Using this option, you use one of the pre-defined services (common protocols such as HTTPS, HTTP, and others) as the protocol to match for the rule to be applied. You can also specify a network service that you configure by navigating to the **Configuration > Roles & Policies > Policies** tab. Select a policy created and click **+** to create a Rule. Select the **Access Control** option in the **Rule Type**. Select the service type from the **Service/app** drop-down list. ■ (other than TCP or UDP) by configuring the IP protocol value.
Action (required)	The action that you want the managed device to perform on a packet that matches the specified criteria. This can be one of the following: ■ permit: Permits traffic matching this rule. ■ drop: Drops packets matching this rule without any notification. ■ reject: Drops the packet and sends an ICMP notification to the traffic source. ■ src-nat: Performs NAT on packets matching the rule. When this option is selected, you need to select a NAT pool. Source IP changes to the outgoing interface IP address (implied NAT pool) or from the pool configured (manual NAT pool). This action functions in tunnel or decrypt-tunnel forwarding mode. ■ dst-nat: This option redirects traffic to the configured IP address and destination port. An example of this option is to redirect all HTTP packets to the captive portal port on the Aruba managed device as used in the pre-defined policy called captive portal. This action functions in tunnel or decrypt-tunnel forwarding mode. User should configure the NAT pool in the managed device. ■ dual-nat: This option performs both source and destination NAT on packets matching the rule. Forward packets from source network to destination; re-mark them with destination IP of the target network. This action functions in tunnel or decrypt-tunnel forwarding mode. User should configure the NAT pool in the managed device. ■ redirect to tunnel: This option redirects traffic into a GRE tunnel. This option is

Parameter	Description
	used primarily to redirect all guest traffic into a GRE tunnel to a DMZ router or switch. ■ redirect to esi: This option redirects traffic to the specified ESI group. You also specify the direction of traffic to be redirected: forward, reverse, or both directions. Select a NAT Pool from the **NAT Pool** drop-down list to add a NAT-POOL for ESI policy. ■ route: Specify the next hop to which packets are routed, which can be one of the following: • Forward Regularly: Packets are forwarded to their next destination without any changes. • Forward to ipsec-map: Packets are forwarded through an IPsec tunnel defined by the specified IPsec map. • Forward to next-hop-list: packets are forwarded to the highest priority active device on the selected next hop list. For more information on next-hop lists, see Uplink Routing using Next-hop Lists on page 255. • Forward to tunnel: Packets are forwarded through the tunnel with the specified tunnel ID. For more information on GRE tunnels, see GRE Tunnels on page 100. • Forward to tunnel group: Packets are forwarded through the active tunnel in a GRE tunnel group. For more information on tunnel groups, see GRE Tunnel Groups on page 107.
TOS (optional)	Value of TOS bits to be marked in the IP header of a packet matching this rule when it leaves the managed device.
Time Range	You can create an **absolute** time range with a single fixed start and end date and time, or create a **periodic** (recurring) time range that starts and ends at a specified time on a weekday, weekend, or selected day.
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. To configure time range, navigate to **Configuration > Roles & Policies > Roles** tab. Select a role and click + in the **Global Rules** table. Select a time range from the **Time range** drop-down list.
Pause ARM Scanning (optional)	Pause ARM scanning while traffic is present. Note that you must enable VoIP Aware Scanning in the ARM profile for this feature to work.
Black List (optional)	Automatically blacklists a client that is the source or destination of traffic matching this rule. This option is recommended for rules that indicate a security breach where the blacklisting option can be used to prevent access to clients that are attempting to breach the security.

Parameter	Description
ACL White List (optional)	A rule must explicitly permit a traffic session before it is forwarded to the managed device. The last rule in the white list denies everything else. Configure white list ACLs on the **Configuration > Services > Firewall > ACL White List** accordion.
802.1p Priority (optional)	When this parameter is enabled, the value of 802.1p priority bits are marked in the frame of a packet matching this rule when it leaves the managed device. 0 is the lowest priority (background traffic) and 7 is the highest (network control).

Creating a Network Service Alias

When you create a network service alias, you can use that alias when specifying the network service for multiple session ACLs.

The following procedure describes how to create a network service alias:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Policies** tab.
2. Click **+** to create a new policy.
3. Enter the policy name in the **Policy name** field.
4. Select a the policy type from the **Policy type** drop-down list. You can select **Ethertype**, **Extended,MAC, Route, Session**, or **Standard**.
5. Click **Submit**.
6. Select the policy created and click **+** in the **Policy <policy name>** table.
7. Select **Access Control** option in the **Rule Type** field.
8. Click **OK**.
9. Select **Service** from the **Service/app** drop-down list.
10. Click + in the **Service alias** drop-down list to add a new service.
    a. Enter a **Service name** .
    b. In the **Protocol** drop-down, select either **TCP** or **UDP**, or select **protocol** and enter the IP protocol number and select an **Application level gateway (alg)** of the protocol for which you want to create an alias.
    c. In the **Port type** drop-down, specify whether you want to define the port by a contiguous range of ports, or by a list of non-contiguous port numbers.
        - If you select **range**, enter the starting and ending port numbers in the **Starting port** and **End port** fields.
        - If you select **list**, enter a comma-separated list of port numbers in the **Port list** field.
    d. To limit the service alias to a specific application, select one the of the following service types from the **Application Level Gateway (alg)** drop-down list:
        - ftp: Service is FTP
        - tftp: Service is TFTP
        - dns: Service is DNS
        - dhcp: Service is DHCP
        - sip: Service is SIP
        - sips: Service is Secure SIP
        - svp: Service is SVP

- sccp: Service is SCCP
- rtsp: Service is RTSP
- vocera: Service is VOCERA
- noe: Service is Alcatel NOE
- h323: Service is H323
- jabber: Service is Jabber
- facetime: Service is Facetime

11. Click **Submit** to add a new service.
12. Click **Submit**.
13. Click **Pending Changes**.
14. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command defines a service alias:

```
(host) [md] (config) #netservice <name> <protocol>|tcp|udp {list <port>,<port>}|{<port>
[<port>]}[ALG <service>]
```

Creating an ACL White List

The white list protects the managed device during traffic session processing by prohibiting traffic from being automatically forwarded to the managed device if it was not specifically denied in a blacklist. The maximum number of entries allowed in the ACL white list is 256. To create an ACL white list, you must first define a white list bandwidth contract, and then assign it to an ACL.

Creating a Bandwidth Contract

The following procedure describes how to create a bandwidth contract:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > Firewall** tab.
2. Expand the **White List BW Contracts** accordion.
3. Click **+** to create a new contract.
4. In the **White list contract name** field, enter the name of a bandwidth contract.
5. In the **Bandwidth rate** field, enter a bandwidth rate value.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command creates a bandwidth contract:

```
(host) [mynode] (config) #cp-bandwidth-contract
```

Configuring the ACL White List

The following procedure describes how to configure an ACL white list:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > Firewall** tab.
2. Expand the **Acl White List** accordion.
3. Click **+** to create a new protocol.
4. Select **permit** or **deny** from the **Action** drop-down list.

   Permit allows session traffic to be forwarded to the managed device and deny blocks session traffic.
5. Select **Ipv4** or **Ipv6** filter from the **IP version** drop-down list.

6. Select one of the following from the **Source** drop-down list:
   - For a specific IPv4 or IPv6 filter, select **addr_mask**. Enter the IP address and mask of the IPv4 or IPv6 filter in the corresponding fields.
   - For a IPv4 or IPv6 host, select **any**.
7. Enter the **IP address** and Subnet **Mask**.
8. In the **IP protocol number(1-255) or IP protocol** field, enter the number for a protocol and select the protocol from the drop-down list used by session traffic.
9. 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.
10. In the **End port** 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.
11. (Optional) Select the name of the bandwidth contract to which the session traffic should be applied, from the **White list bandwidth contract** drop-down list.
12. For further information on creating bandwidth contracts, see Using an Exclude List on page 520
13. Click **Submit**. The ACL displays on the white list section.
14. To delete an entry, click **Delete** next to the entry you want to delete.
15. Click **Submit**.
16. Click **Pending Changes**.
17. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command creates ACL white lists:

```
(host) [mynode] (config)firewall cp
```

Override Local Network Destination

To implement this feature, a new sub-command, **host vlan – offset** under the **netdestination** configuration command is introduced. An example and description are as follows:

```
netdestination store
   host vlan 10 offset 5
   host vlan 10 offset 8
```

With the above, select the subnet (for example, 10.1.1.0/24) assigned to vlan 10 for that store and calculate offsets 5 (10.1.1.5) and 8 (10.1.1.8) from it.

The following procedure describes how to configure an override local network destination:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Roles** tab.
2. Select a role and click **+** under **Rules of this Role only** to create a rule.
3. Click one of the options in the **Rule Type** filed to select a rule and click **OK**.
4. Select **Alias** from the **Destination** drop-down list.
5. Select **+** from the **Destination alias** drop-down list.
6. Click **+** in the **Rule** table.
7. Select **Override** from the **Rule type** drop-down list.
8. Select a VLAN offset number which is the Netmask or range, from the **Vlan** drop-down list.
9. Click **OK**.
10. Click **Submit** in the **Add New Destination** window.
11. Click **Submit**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the local override netdestination:

```
(host) [md] (config) #netdestination store
(host) [md] (config-submode) #?
  description             Brief description about this destination (up to 128 characters
  in quote)
  host                    Configure a single IPv4 host
  invert                  Use all destinations EXCEPT this destination
  name                    Configure a single host name or domain, Max 63 characters
  network                 Configure a IPv4 subnet
  no                      Delete Command
  range                   Configure a range of IPv4 addresses
(host) [md] (config-submode) #host?
  vlan                    IPv4 Address based on VLAN
  A.B.C.D                 IPv4 Address of host
(host) [md] (config-submode) #host vlan ?
  <1-4094>                VLAN ID
(host) [md] (config-submode) #host vlan 55 ?
  offset                  Offset in the VLAN subnet
(host) [md] (config-submode) #host vlan 55 offset ?
  <1-254>                 Offset number in the VLAN subnet
(host) [md] (config-submode) #host vlan 55 offset 36
```

Execute the following command to show the local override netdestination:

```
(host) [md] #show netdestination store
  Name: store
  Position  Type      IP addr  Mask-Len/Range
  --------  ----      -------  --------------
  1         override  vlan 55  offset 36
```



One netdestination definition can have a maximum of 256 netdestination entries. On the whole, there can be a maximum of 1024 netdestination entries on the Controller or Managed Device.

How to use the local-override netdestination alias in the managed device:

```
(host) [md] (config) #ip access-list  session store-override
(host) [md] (config-sess-store-override) #any alias store any permit
(host) [md] (config-sess-store-override) #alias store any any deny
(host) [md] (config-sess-store-override) #!
(host) [md] #show ip interface brief
  Interface                  IP Address / IP Netmask        Admin   Protocol
  vlan 1                  172.72.10.254 / 255.255.255.0     up      up
  vlan 55                    55.55.55.1 / 255.255.255.0     up      up
  loopback                   unassigned / unassigned        up      up

(host) [md] #show acl acl-table | include dummy-acl
75   session      620        2           3          dummy-acl                    0

(host) [md] #show acl ace-table acl 75

620: any netdest-id: 34  0  0-0  0-0  f1000080001:permit  alias-dst  hits-table-index
24578
621: netdest-id: 34 any  0  0-0  0-0  f800080001:permit  alias-src  hits-table-index
24579
622: any any  0  0-0  0-0  f180000:deny
```

RTP Traffic without Changing DSCP value

The RTP traffic can be passed without changing the DSCP value set by the end user device. This allows the RTP traffic to pass through the managed devices.

To pass the RTP traffic without changing the DSCP value, execute the following command:

```
(host) [md] (config) #firewall
```

```
(host) [md] (config-submode)#voip-qos-trusted
```

To verify if the RTP traffic is passed without changing the DSCP value, execute the following command:

```
(host) [md] #show firewall | include Trust

Trust packet QoS                                Enabled
```

To verify the client DSCP value (for example, 48) for RTP traffic, execute the following command:

```
(host) #show datapath session dpi | include V

C - client, M - mirror, V - VOIP
r - Route Nexthop, h - High Value

Source IP or MAC  Destination IP  Prot SPort DPort Cntr     Prio ToS Age Destination TAge
10.15.123.147     10.15.16.19     17   33262 2060   0/0     6    48  0   local       2876
10.15.16.19       10.15.123.147   17   2060  33262  0/0     6    48  0   local       2876

Packets    Bytes      AclVer   Int-Flag Sess-Flag2 PktsDpi   UplnkVlan AppID
1          40         8009     81095    0          3         none      alg-rtp
0          0          0        1094     0          2         none      alg-rtp

AceIdx             Flags       DpiTIdx    CPU ID
(3404) 1142/1138   FHPTCVBO    dc         7
(3404) 0/1138      FHPTCVBO    dc         6
```

Creating a User Role

User roles comprises of user role settings, firewall policies, and bandwidth contracts. This section describes the procedure to create and delete a user role, and associate a firewall policy with that role.

The commands to associate an ACL to a user role vary, depending upon the type of ACL being associated to that role. User roles are applied globally across all managed devices, so ethertype, MAC and session ACLs can be applied to global user roles. However, routing access lists may vary between locations, so they are mapped to a user role in a local configuration setting.

- To associate a user role with an ethertype, MAC or session ACL, use the **user-role <role> access-list eth|mac|session <acl>** command.
- To associate a user role with an routing ACL, use the **routing-policy-map** command.

The following procedure describes how to create a new user role:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Roles & Policies** > **Roles** .
2. Click **+** to create a new role.
3. Enter a **Name** for the new role and click **Submit**.
4. Select the role created and click **+** under **Rules of this role only** table.
5. Click one of the options in the **Rule Type** filed to select a rule and click **OK.**
6. In the **New Forwarding Rule** section, configure all the parameters.
7. Click **Submit**.
8. Select one of the following options to add a policy to the role:
   - In the **Policies** tab select the role created and click **+** under the **Policies** table. Enter a **Name** for the policy and select a **Policy type**. Click **Submit**.
   - To associate an existing policy to a user role:
     - Select the **Role** from the **Roles** tab and click **Show Advanced View** in **Roles <policy name>** table.

- Click **+** under the **Policies** tab.
- Select **Add an existing policy** option and select a policy from the **Policy name** drop-down list.
- Click **Submit**.



For more information on creating a firewall policy, see Firewall Policies on page 491.

9. (Optional) If the user role contains more than one firewall policy, use the up and down arrows to assign priorities to each role. The higher the policy on the list, the higher its priority.
10. Click **Show Advanced View** and enter the configuration values as described in Table 82.
11. Click **Submit.**
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.
14. Assign the user role to a AAA profile in the managed device. After assigning the user role, execute the **show reference user-role <role>** command on the managed device to see the profiles that reference this role. For more information, see Workflow for Assigning a User Role on page 504

Table 82: *User Role Parameters*

Parameter	Description
Name	Name of the user role. The character length of a user role is from 1-63 characters.
More	
VLAN (optional)	Navigate to **More > Network** to assign VLAN ID to the user role. By default, a client is assigned a VLAN on the basis of the ingress VLAN for the client to the managed device. You can override this assignment and configure the VLAN ID that is to be assigned to the user role.
Re-auth interval (optional)	Navigate to **More > Network** to configure time, in minutes, after which the client is required to reauthenticate. Enter a value between 0-4096. 0 disables reauthentication. Default: 0 (disabled)
Max Sessions (optional)	Navigate to **More > Network** to configure the maximum number of sessions per user in this role. If the sessions reach the maximum value, any additional sessions from this user that are reaching the threshold are blocked till the session usage count for the user falls back below the configured limit. The default is 65535. You can configure any value between 0-65535.
Deep packet inspection (optional)	Navigate to **More > Network** to enable or disable deep packet inspection. This setting is enabled by default.
Web content classification (optional)	Navigate to **More > Network** to enable or disable web content classification for all HTTP traffic. This setting is enabled by default.
YouTube education (optional)	Navigate to **More > Network** to enable or disable YouTube education. This setting is disabled by default. If enabled, the page redirects to YouTube education where non-educational videos are not streamed and the user can enter a YouTube education enabled cookie (optional).

Parameter	Description
Open flow (optional)	Navigate to **More > Network** to enable or disable Software Defined Network for the user role. This setting is enabled by default.
VPN Dialer (optional)	Navigate to **More > VPN** to assign a VPN dialer to a user role. For details about VPN dialer, see Virtual Private Networks on page 459. 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)	Navigate to **More > VPN** to assign an L2TP pool to the user role. For more details about L2TP pools, see Virtual Private Networks on page 459. 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)	Navigate to **More > VPN** to assign a PPTP pool to the user role. For more details about PPTP pools, see Virtual Private Networks on page 459. 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.
VIA connection profile	Navigate to **More > VPN** to assign a VIA connection profile to the user role.
IDP profile (optional)	Navigate to **More > Authentication** to assign a IDP profile to the user role. For more details, refer to
Stateful NTLM profile (optional)	Navigate to **More > Authentication** to assign a stateful NTLM profile to the user role. For more details, refer to Configuring Stateful NT LAN Manager Authentication on page 313.
Stateful Kerberos profile (optional)	Navigate to **More > Authentication** to assign a stateful Kerberos profile to the user role. For more details, refer to Configuring Stateful Kerberos Authentication on page 315.
WISPr profile (optional)	Navigate to **More > Authentication** to assign a WISPr profile to the user role.For more details, refer to WISPr Authentication on page 312.
Captive Portal Profile (optional)	Navigate to **More > Authentication** to assign a Captive Portal profile to this role. For more details about Captive Portal profiles, see Captive Portal Authentication on page 327.
Captive Portal Check for Accounting (optional)	Navigate to **More > Authentication** to enable or disable this setting. This setting is enabled by default. If disabled, RADIUS accounting is done for an authenticated users irrespective of the captive-portal profile in the role of an authenticated user. If enabled, accounting is not done as long as the user's role has a captive portal profile on it. Accounting will start when Auth or XML-Add or CoA changes the role of an authenticated user to a role which doesn't have captive portal profile.
Bandwidth	
Bandwidth (optional)	Navigate to **Show Advanced View > Bandwidth** to assign a bandwidth contract and 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 Using an Exclude List on page 520.
Captive Portal	
Captive Portal	This tab allows you to personalize the captive portal page. For details, refer to Personalizing the Captive Portal Page on page 353.

Deleting a User Role

The following procedure describes how to delete a user role:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Roles** tab on the WebUI.
2. Select the **Role** and click the **Delete** icon.



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.

Workflow for Assigning a User Role

A client is assigned a user role by one of several methods. A role assigned by one method may take precedence over one assigned by a different method.

The methods of assigning user roles are mentioned below. The tasks are set in the precedence of lowest to highest.

1. The initial user role or VLAN for unauthenticated clients is configured in the AAA profile for a virtual AP.

   See Access Points on page 627 and Assigning User Roles in AAA Profiles on page 504.
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*. UDRs are executed *before* client authentication.

   See Working with User-Derived VLANs on page 505.
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.

   See Configuring a Default Role for Authentication Method on page 508.
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.

   See Configuring a Server-Derived Role on page 509.
5. The user role can be derived from Aruba VSA for RADIUS server authentication. A role derived from an Aruba VSA takes precedence over any other user roles.

   See Configuring a VSA-Derived Role on page 509.

Assigning User Roles in AAA Profiles

An AAA profile defines the user role for unauthenticated clients (initial role) as well as the default user role for MAC and 802.1X authentication. For additional information on creating AAA profiles, see WLAN Authentication on page 568.

The following procedure describes how to assign user roles in AAA profiles:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > AAA Profiles** tab.
2. Expand the **AAA Profiles** and select a AAA profile.
3. Select the default profile or a user-defined AAA profile.
4. Select the desired user role for unauthenticated users, from the **Initial Role** drop-down list.
5. Select the desired user role for users who have completed 802.1X authentication, from the **802.1X Authentication Default Role** drop-down list.
6. Select the desired user role for clients who have completed MAC authentication, from the **MAC Authentication Default Role** drop-down list.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures user roles in AAA profile:

```
(host) [md] (config) #aaa profile <profile-name>
```

Working with User-Derived VLANs

Attributes derived from the client's association with an AP can be used to assign the client to a specific role or VLAN, as UDRs are executed before the client is authenticated.

You configure the user role or VLAN to be assigned to the client by specifying condition rules; when a condition is met, the specified user role or VLAN is assigned to the client. You can specify more than one condition rule; the order of rules is important as the first matching condition is applied. You can optionally add a description of the user rule.

The Table 83 describes the conditions for which you can specify a user role or VLAN.

Table 83: *Conditions for a User-Derived Role or VLAN*

Rule Type	Condition	Value
BSSID: Assign client to a role or VLAN based upon the BSSID of AP to which client is associating.	One of the following: ■ contains ■ ends with ■ equals ■ does not equal ■ starts with	MAC address (xx:xx:xx:xx:xx:xx)
DHCP-Option: Assign client to a role or VLAN based upon the DHCP signature ID.	One of the following: ■ equals ■ starts with	DHCP signature ID. **NOTE:** This string is *not* case sensitive.
DHCP-Option-77: Assign client to a role or VLAN based upon the user class identifier returned by DHCP server.	equals	string
Encryption: Assign client to a role or VLAN based upon the encryption type used by the client.	One of the following: ■ equals ■ does not equal	■ Open System (no encryption) ■ WPA or WPA2 AES (static or dynamic) ■ WPA or WPA2-TKIP (static or dynamic)

Rule Type	Condition	Value
		■ WEP (static or dynamic) ■ xSec
ESSID: Assign client to a role or VLAN based upon the ESSID to which the client is associated.	One of the following: ■ contains ■ ends with ■ equals ■ does not equal ■ starts with ■ value of (does not take *string*; attribute value is used as role)	string
Location: Assign client to a role or VLAN based upon the AP name to which the client is associated.	One of the following: ■ equals ■ does not equal	string
MAC address of the client	One of the following: ■ contains ■ ends with ■ equals ■ does not equal ■ starts with	MAC address (xx:xx:xx:xx:xx:xx)

Understanding Device Identification

The device identification feature allows you to assign a user role or VLAN to a specific device type by identifying a DHCP option and signature for that device. If you create a user rule with the **DHCP-Option** rule type, the first two characters in the **Value** field must represent the hexadecimal value of the DHCP option that this rule should match, while the rest of the characters in the **Value** field indicate the DHCP signature the rule should match. To create a rule that matches DHCP option 12 (host name), the first two characters in the **Value** field must be the hexadecimal value of 12, which is 0C. To create a rule that matches DHCP option 55, the first two characters in the Value field must be the hexadecimal value of 55, which is 37.

The following table describes some of the DHCP options that are useful for assigning a user role or VLAN.

Table 84: *DHCP Option values*

DHCP Option	Description	Hexadecimal Equivalent
12	Host name	0C
55	Parameter Request List	37
60	Vendor Class Identifier	3C
81	Client FQDN	51

The device identification features in ArubaOS can also automatically identify different client device types and operating systems by parsing the User-Agent strings in the client's HTTP packets. To enable this feature, select the **Device Type Classification** option in the AP's AAA profile. For details, see WLAN Authentication on page 568.

Starting from ArubaOS 8.0.1, the device type classification is enhanced to identify the device type for each client, determine firewall policies, and customize to meet the requirement of the end user. The device type information is sent from ClearPass to ArubaOS.



Prior to establishing the WebSocket interface with ClearPass Insight server the issuer certificate of the server must be imported to the controller as TrustedCA certificate.

To gather the information required to manage and establish WebSocket interface to the ClearPass Insight server, configure ClearPass WebSocket profile. Once the connection is established, the user can subscribe or unsubscribe and receive device profile information for the subscribed stations.



Only admin, apiadmin, and clusteradmin can configure ClearPass WebSocket profile.

The following procedure describes how to configure the ClearPass WebSocket interface and the primary and secondary ClearPass Insight server:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From **All Profiles** select **Other Profiles > ClearPass WebSocket**.
3. Select ClearPass WebSocket Interface checkbox to enable this option and to connect to ClearPass WebSocket.
4. Enter appropriate values in the **host** and **portnum** fields.
5. Enter appropriate values in the parameters listed below the **Primary ClearPass Insight Server** and **Secondary ClearPass Insight Server** fields.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the ClearPass WebSocket interface and the primary and secondary ClearPass Insight server:

```
(host) [mynode] (config) #websocket clearpass
(host) [mynode] (ClearPass WebSocket Profile) #primary host <host> port <1-65535>
username <username> passwd <passwd>
(host) [mynode] (ClearPass WebSocket Profile) #secondary host <host> port <1-65535>
username <username> passwd <passwd>
(host) [mynode] (ClearPass WebSocket Profile) #enable
```

The following CLI command checks the current connection state of the ClearPass WebSocket interface:

```
(host) [mynode] #show websocket state clearpass
```

The following CLI command helps to view the current statistics of ClearPass WebSocket interface:

```
(host) [mynode] #show websocket statistics clearpass
```

Configuring a User-derived VLAN

The following procedure describes how to configure a user derived VLAN:

1. In the **Managed Device** node hierarchy, navigate to **Configuration** > **Authentication** > **User Rules** tab.

2. Click **+** to add a new set of derivation rules. Enter a **Name** for the set of rules, and click **Submit**.
   The name appears in the **User Rules Summary** list.
3. In the **User Rules Summary** list, select the name of the rule created to configure rules.
4. Click **+** in the **Rules-set** table to add a rule.
5. Select **VLAN** from the **Set Type** drop-down list.
   You can select **VLAN** to create derivation rules for setting the VLAN assigned to a client.
6. Configure the condition for the rule by setting the **Rule type**, **Condition**, **Value** parameters and optional description of the rule. See Table 83 for descriptions of these parameters.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.
10. 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.)
11. (Optional) If the rule uses the DHCP-Option condition, the best practice is to enable the **Enforce DHCP** parameter in the AP group's AAA profile, which requires users to complete a DHCP exchange to obtain an IP address. For details on configuring this parameter in an AAA profile, see WLAN Authentication on page 568.

When you create a user derivation rule by selecting **VLAN** from the **Set Type** drop-down list, you must configure the AP group's AAA profile to use the rule. For more information, see WLAN Authentication on page 568

The following CLI command configures a user derived VLAN:

```
(host) [md] (config) #aaa derivation-rules user <name>
```

The following CLI commands configure a AAA profile with user derivation rule:

```
(host) [md] (config) #aaa profile <profile_name>
(host) [md] (AAA Profile <profile_name>) #user-derivation-rules <rule_name>
```

RADIUS Override of User-Derived Roles

This feature introduces a new RADIUS vendor specific attribute (VSA) named Aruba-No-DHCP-Fingerprint, value 14. This attribute signals the RADIUS Client (managed device) to ignore the DHCP Fingerprint user role and VLAN change post L2 authentication. This feature applies to both Campus AP and Remote AP in tunnel forwarding mode and for the L2 authenticated role only.

Configuring a Default Role for Authentication Method

For each authentication method, you can configure a default role for clients who are successfully authenticated using that method.

The following procedure describes how to configure a default role for an authentication:

1. In the **Managed Network** node hierarchy, navigate to **Configuration** > **Authentication** > **AAA Profiles** tab.
2. To configure the default user role for MAC or 802.1X authentication, select a AAA profile under **AAA Profiles** and select the desired 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.

a. For L2 Authentication, select Stateful 802.1X authentication type and select the user role for **Default role**.

b. For L3 Authentication, select the authentication type (Captive Portal or VPN Authentication) and then select a profile. Select the user role for **Default Role**.

4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

For additional information on configuring captive portal authentication, see Captive Portal Authentication on page 327.

The following CLI command configures the default user role for MAC or 802.1X authentication:

```
(host) [md] (config) #aaa profile <profile>
```

The following CLI command configures the default user role for other authentication methods:

```
(host) [md] (config) #aaa authentication captive-portal|stateful-dot1x|stateful-ntlm|vpn
```

Configuring a Server-Derived Role

If the client is authenticated through an authentication server, the user role for the client can be based on one or more attributes returned by the server during authentication. You configure the user role to be derived by specifying condition rules; when a condition is met, the specified user role is assigned to the client. You can specify more than one condition rule; the order of rules is important as the first matching condition is applied. You can also define server rules based on client attributes such as ESSID, BSSID, or MAC address, even though these attributes are not returned by the server.

For information about configuring a server-derived role, see Configuring Server-Derivation Rules on page 212.

Configuring a VSA-Derived Role

Many Network Address Server (NAS) vendors, including Aruba, use VSAs to provide features not supported in standard RADIUS attributes. For Aruba 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 (Aruba) 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 managed devices conform to the format recommended in RFC 2865, “Remote Authentication Dial In User Service (RADIUS)”.

For more information on Aruba VSAs, see Configuring Authentication Servers on page 193. Dictionary files that contain Aruba VSAs are available on the Aruba support website for various RADIUS servers. Log into the Aruba support website to download a dictionary file from the Tools folder.

Understanding Global Firewall Parameters

Each firewall policy has a each of parameters that require configuration. In order to set up robust firewall policies, it is essential to understand what each parameter does, it's functionality, and purpose. Table 85 describes optional firewall parameters you can set on the managed devices for IPv4 traffic.

To configure global firewall parameters, in the **Mobility Master** node hierarchy, navigate to the **Configuration** > **Services** > **Firewall** > **Global Settings** accordion and select or enter values in the IPv4 column.

You can also use the CLI command **firewall** for configuration.

See IPv6 Support on page 122 for information about configuring firewall parameters for IPv6 traffic.

Table 85: *Pv4 Firewall Parameters*

Parameter	Description
Monitor Ping Attack (per 30 seconds)	Number of ICMP pings per 30 second, which if exceeded, can indicate a DoS attack. Valid range is 1-16384 pings per 30 seconds. Recommended value is 120 packets per 30 seconds. Default: No default
Monitor TCP SYN Attack rate (per 30 seconds)	Number of TCP SYN messages per 30 second, which if exceeded, can indicate a DoS attack. Valid range is 1-16384 pings per 30 seconds. Recommended value is 960 packets per 30 seconds. Default: No default
Monitor IP Session Attack (per 30 seconds)	Number of TCP or UDP connection requests per 30 second, which if exceeded, can indicate a DoS attack. Valid range is 1-16384 requests per 30 seconds. Recommended value is 960 packets per 30 seconds. Default: No default
Monitor/Police ARP Attack (non Gratuitous ARP) rate (per 30 seconds)	Number of ARP packets (other than Gratuitous ARP packets) per 30 seconds, which if exceeded, can indicate a DoS attack. Valid range is 1-16384 packets per 30 seconds. Recommended value is 960 packets per 30 seconds. Default: No default
Monitor/Police Gratuitous ARP Attack rate (per 30 seconds)	Number of Gratuitous ARP packets per 30 seconds, which if exceeded, can indicate DoS attack. Valid range is 1-16384 packets per 30 seconds. Recommended value is 50 packets per 30 seconds. Default: 50 packets
Monitor/Police Gratuitous ARP Attack Action	Select **Blacklist** to block the gratuitous ARP or **Drop** to disallow a gratuitous ARP from untrusted ports.
Monitor/Police CP Attack rate (per 30 seconds)	Rate of misbehaving user's traffic, which if exceeded, can indicate a denial or service attack. Recommended value is 3000 frames per 30 seconds. Default: No default
Deny Inter User Bridging	Prevents the forwarding of Layer-2 traffic between wired or wireless users. You can configure user role policies that prevent Layer-3 traffic between users or networks but this does not block Layer-2 traffic. This option can be used to prevent traffic, such as Appletalk or IPX, from being forwarded. Default: Disabled
Deny Inter User Traffic	Denies traffic between untrusted users by disallowing layer-2 and layer-3 traffic. This parameter does not depend on the deny-inter-user-bridging parameter being enabled or disabled. Default: Disabled
Deny Source Routing	Permits the firewall to reject and log packets with the specified IP options loose source routing, strict source routing, and record route. Note that network packets where the IPv6 source or destination address of the network packet is defined as an link-local address (fe80::/64) are permitted. Default: Disabled
Deny All IP Fragments	Drops all IP fragments.

Parameter	Description
	NOTE: Do not enable this option unless instructed to do so by an Aruba representative. Default: Disabled
Enforce TCP Handshake Before Allowing Data	Prevents data from passing between two clients until the three-way TCP handshake has been performed. This option should be disabled when you have mobile clients on the network as enabling this option will cause mobility to fail. You can enable this option if there are no mobile clients on the network. Default: Disabled
Prohibit IP Spoofing	Enables detection of IP spoofing (where an intruder sends messages using the IP address of a trusted client). When this option is enabled, source and destination IP and MAC addresses are checked for each ARP request or response. Traffic from a second MAC address using a specific IP address is denied, and the entry is not added to the user table. Possible IP spoofing attacks are logged and an SNMP trap is sent. Default: Enabled
Prohibit RST Replay Attack	When enabled, closes a TCP connection in both directions if a TCP RST is received from either direction. You should not enable this option unless instructed to do so by an Aruba representative. Default: Disabled
Log all received ICMP Errors	Enables logging of received ICMP errors. You should not enable this option unless instructed to do so by an Aruba representative. Default: Disabled
Stateful SIP Processing	Disables monitoring of exchanges between a VoIP or VoWLAN 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 managed device is *not* the default gateway for wireless clients and the default gateway is behind the managed device. This option is typically used for captive portal configuration. Default: Disabled.
AMSDU Configuration	Enables handling AMSDU traffic from clients. Default: Disabled
Session Idle Timeout (sec)	Set the time, in seconds, that a non-TCP session can be idle before it is removed from the session table. Specify a value in the range 16-300 seconds. You should not set this option unless instructed to do so by an Aruba representative. Default: 16 seconds
Session Mirror Destination	Destination (IP address or port) to which mirrored session packets are sent. This option is used only for troubleshooting or debugging. Packets can be mirrored in multiple ACLs, so only a single copy is mirrored if there is a match within more than one ACL. You can configure the following: ▪ Ethertype to be mirrored with the Ethertype ACL mirror option. ▪ IP flows to be mirrored with the session ACL mirror option.

Parameter	Description
	▪ MAC flows to be mirrored with the MAC ACL mirror option. ▪ If you configure both an IP address and a port to receive mirrored packets, the IP address takes precedence. Default: N/A
Disable FTP Server	Disables the FTP server on the managed device. Enabling this option prevents FTP transfers. You should not enable this option unless instructed to do so by an Aruba representative. Default: Disabled (FTP server is enabled)
GRE Call ID Processing	Creates a unique state for each PPTP tunnel. You should not enable this option unless instructed to do so by an Aruba representative. Default: Disabled
Optimize duplicate access detection frames	Optimizes DAD frames and reduces flooding of IPv4 gARPs / IPv6 DAD frames onto wireless clients. Default: Enabled
Stall detection	Triggers datapath crash on stall detection. This is applicable only to 7200 Series controllers. You should not enable this option unless instructed to do so by an Aruba representative. Default: Disabled
Immediate freeback	If enabled, it immediately frees buffers on controllers. You should not enable this option unless instructed to do so by Aruba representative.
Stateful ICMP processing	It creates sessions for ICMP errors and denies unidirectional replies. Default: Disabled
Mcast RED	Configures multicast random early detection algorithms. Click the toggle switch to enable this setting. The following parameters are displayed only when **Mcast RED** is enabled: ▪ **Inverse mark probability**—Specify an Inverse mark probability value. For example, an inverse mark probability parameter of 10 corresponds to a mark probability of 1/10 which means 1 in 10 packets will be dropped. ▪ **Minimum threshold**—Specify a minimum threshold value. Range is 0-99. ▪ **Maximum threshold**—Specify a maximum threshold value. Range is 1-100.
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 Aruba representative, as doing so may create unnecessary overhead on the managed device. Default: Disabled (per-session logging is performed)
Broadcast-filter ARP	Reduces the number of broadcast packets sent to VoIP clients, thereby improving the battery life of voice handsets. You can enable this option for voice handsets in conjunction with increasing the DTIM interval on clients. Default: Disabled

Parameter	Description
Prohibit ARP Spoofing	Detects and prohibits ARP spoofing. When this option is enabled, possible arp spoofing attacks are logged and an SNMP trap is sent. Default: Disabled
Prevent DHCP Exhaustion	Enable check for DHCP client hardware address against the packet source MAC address. This command checks the frame's source-MAC against the DHCPv4 client hardware address and drops the packet if it does not match. Enabling this feature prevents a client from submitting multiple DHCP requests with different hardware addresses, thereby preventing DHCP pool depletion. Default: Disabled
Only Allow Local Subnets in User Table	Adds only IP addresses, which belong to a local subnet, to the user-table. Default: Disabled
Session-tunnel FIB	Enable session-tunnel based forwarding. **NOTE:** Best practices is to enable this parameter only during maintenance window or off-peak production hours.
Multicast Automatic Shaping	Enables multicast optimization and provides excellent streaming quality regardless of the amount of VLANs or IP IGMP groups that are used. Default: Disabled
Enforce BW Contracts for Broadcast Traffic	Applies bw contracts to local subnet broadcast traffic.
Enforce TCP Sequence Numbers	Enforces the TCP sequence numbers for all packets. Default: Disabled
Session VOIP Timeout (sec)	Sets the idle session timeout for sessions that are marked as voice sessions. If no voice packet exchange occurs over a voice session for the specified time, the voice session is removed. Range is 16 – 300 seconds. Default: 300 seconds
Stateful H.323 Processing	Disables stateful H.323 processing. Default: Enabled
Stateful SCCP Processing	Disables stateful SCCP processing. Default: Disabled
Session Mirror IPSEC	Configures session mirroring of all frames that are processed by IPsec. Frames are sent to IP address specified by the session-mirror-destination option. **NOTE:** Use this option for debugging or troubleshooting only. Default: Disabled
Stateful VOCERA Processing	Disables stateful VOCERA processing. Default: Disabled
Stateful UA Processing	Disables stateful UA processing. Default: Disabled

Parameter	Description
Enforce WMM Voice Priority Matches Flow Content	If traffic to or from the user is inconsistent with the associated QoS policy for voice, the traffic is reclassified to best effort and data path counters incremented. Default: Disabled
Rate Limit CP Untrusted Ucast Traffic (pps)	Specifies the untrusted unicast traffic rate limit. Range is 1-65535 packets per seconds (pps). Default: 9765 pps
Rate Limit CP Untrusted Mcast Traffic (pps)	Specifies the untrusted multicast traffic rate limit. Range is 1-65535 packets per seconds (pps). Default: 1953 pps
Rate Limit CP Trusted Ucast Traffic (pps)	Specifies the trusted unicast traffic rate limit. Range is 1-98304 packets per seconds (pps). Default: Disabled
Rate Limit CP Trusted Mcast Traffic (pps)	Specifies the trusted multicast traffic rate limit. Range is 1-65535 packets per seconds (pps). Default: 1953 pps
Rate Limit CP Route Traffic (pps)	Specifies the traffic rate limit that needs ARP requests. Range is 1-65535 packets per seconds (pps). Default: 976 pps
Rate Limit CP Session Mirror Traffic (pps)	Specifies the session mirrored traffic forwarded to the managed device. Range is 1-65535 packets per seconds (pps). Default: 976 pps
Rate limit CP VRRP traffic (pps)	Rate of the VRRP traffic hitting the control plane. Default: 512 pps
Rate limit CP ARP traffic(pps)	Rate of the ARP traffic hitting the control plane. Default: 976 pps
Rate limit CP I2 protocol / other traffic (pps)	Rate of other L2 traffic (non- IP and ARP) hitting the control plane.
Rate Limit CP Auth Process Traffic (pps)	Specifies the traffic rate limit that is forwarded to the authentication process. Range is Range is 1-65535 packets per seconds (pps). Default: 976 pps
Rate Limit CP IKE Traffic	The bandwidth contract for CP IKE traffic. Default: 1953 pps
Jumbo Frames Processing	Enables jumbo frame processing for data frames that are larger than 1500 bytes. Default: Disabled
Enable deep packet inspection	If enabled, it performs deep packet inspection. Default: Disabled
Enable web content classification	Enables web content classification for all HTTP traffic.
Drop packets using web content cache miss	Drops data packets that do not match any web content category or reputation levels in the

Parameter	Description
	managed device's internal web content cache. Default: Disabled

Working in the Presence of Web Proxy

When the Mobility Master needs to access data on the cloud or the internet, and if the internet bound traffic needs to pass through a proxy, execute the **web-proxy server** command. Once the command is executed the Mobility Master routes web (HTTP or HTTPS) traffic through the proxy server.

Execute the following command in the CLI to route web traffic through the proxy server:

```
(host) [mynode] (config) #web-proxy server arubaproxy.com port 8080
(host) [mynode] (config) #show web-proxy
                Server: arubaproxy.com
                port: 8080
```

Support for Desktop Virtualization Protocols

ArubaOS supports desktop virtualization protocols by providing preconfigured ACLs for Citrix and VMware clients. You can apply these ACLs to the user-role when using the Virtual Desktop Infrastructure clients. This ensures that any enterprise application that uses the VDI client performs optimally with appropriate QoS.



Disable the voice aware ARM when applying the ACLs for the VDI clients as the virtual desktop sessions may prevent the ARM scanning.

Configuring Firewall Settings for Protection from ARP Attacks

The following procedure describes how to configure firewall settings to protect the network against attacks:

1. In the Mobility Master node hierarchy, navigate to **Configuration** > **Services** > **Firewall** tab.
2. Under **Software Management** click **Reboot**.

Blacklisting Wired Clients

Starting ArubaOS 8.2.0.0, you can blacklist wired clients. This feature is useful where firewall policies are applied for wired traffic. For example, remote APs in which wired ports are used or remote APs in tunneled node.

The following CLI command configures the blacklist timer for a wired client:

```
(host) [mynode] (config) #aaa authentication wired
(host) [mynode] (Wired Authentication Profile) # blacklist-time <timer>
```

Limitations

Blacklisting wired clients has certain limitations also. The limitations of this feature are:

- Functions only for wired clients on tunnel-based remote APs for secure jack operation.
- Supports blacklisting wired clients based on number of ACL entry hits.
- Is not supported in a cluster topology.

AppRF 2.0

The AppRF 2.0 feature improves application visibility and control by allowing you to configure ACL and bandwidth-control applications and application categories. AppRF 2.0 supports a DPI engine for application detection for over a thousand applications. All wired and wireless traffic that traverses the managed device can now be categorized and controlled by application and application category.

AppRF 2.0 provides the ability to:

- permit or deny an application or application category for a specific role. For example, you can block bandwidth monopolizing applications on a guest role within an enterprise.
- rate limit an application or application category, such as video streaming applications, globally or for a specific role.
- mark different L2 or L3 QoS for an application or application category for a user role. For example, you can mark video and voice sessions that originate from wireless users with different priorities so that traffic is prioritized accordingly in your network.
- support for upgrading application signatures.
- define custom applications and custom application categories.

Workflow for configuring AppRF 2.0

Configuring AppRF improves application visibility and control. To configure AppRF, perform the following tasks:

1. Enable a DPI configuration on the managed device.
   See Enabling DPI on page 516.
2. Configure a policy to permit or deny an application or application category on a given role.
   See Configuring Policies for AppRF 2.0 on page 517.
3. Configure bandwidth contracts for both the global or application-specific levels.
   See Configuring Bandwidth Contracts on page 519.
4. Upgrade the application signature using protocol based image upgrade.
   See Upgrading Application Signatures on page 521.
5. Define a custom application for users to apply roles and policies, and bandwidth contracts.
   See Defining Custom Application on page 521.

Enabling DPI

For application and application category specific configuration to take affect, you must first enable DPI.



You must reboot (reload) the managed device after you enable or disable DPI for global classification to take effect.

The following procedure describes how to enable DPI:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > Firewall** tab.
2. Expand the **Global Settings** accordion.
3. Select the **Enable deep packet inspection** check box.
4. Click **Submit**.
5. Click **Pending Changes**.

6. In the **Pending Changes** window, select the check box and click **Deploy changes**.
7. Reload the Mobility Master.

The following CLI command enables DPI:

```
(host) [mynode](config) #firewall dpi
(host) [mynode] #reload
```

The following CLI command displays the application ID, application name, and the ACL or ACE index information for a given session:

```
(host) [mynode] #show datapath session dpi
```

Configuring Policies for AppRF 2.0

ACL now contain new application and application category options that let you permit or deny an application or application category on a given role. See the Dashboard Monitoring Traffic Analysis topic for details about configuring policies from the Dashboard.

How ACL Works with AppRF

A session entry proceeds through two phases: the application detection phase (phase 1) and the post-application detection phase (phase 2). A session ACL is applied in phase 1 and in phase 2.

In phase 1, if the session ACL lookup results in an L3 or L4 ACE entry request, the traffic pertaining to the session is guided by this L3 or L4 ACE entry. However, if the session ACL lookup results in an application or application category specific ACE entry, the enforcement is postponed until phase 2. Once the application is determined, the session ACL is re-applied with application or application category information to determine the final action on the traffic.

Global Session ACL

The Global Session ACL is used to configure ACL rules that span across or are common to all roles. They are applied to all roles. The global-sacl rules take precedence over any other ACLs that may be in the user role.

The global-sacl session ACL by default, is in position one for every user role configured on the managed device. The global-sacl session ACL has the following properties:

- It cannot be deleted.
- It always remains at position one in every role and its position cannot be modified.
- It contains only application rules.
- It can be modified in the WebUI, CLI, and dashboard on a Mobility Master.
- Any modifications to it results in the regeneration of ACE's of all roles.

Role Default Session ACL

You can configure role-specific application configuration using the WebUI and dashboard. For example, you can deny the facebook application on the guest role using the CLI or dashboard without having to change the firewall configuration. This per-user role configuration from WebUI or Dashboard is placed in the Role Default Session ACL.

A new role session ACL named apprf-role-name-sacl has been added. This session, by default, is in position two for every user role configured on the managed device.

The string apprf is added to the beginning and sacl to the end of a role's name to form a managed device unique name for role default session ACL. This session ACL is in position two of the given user role after the global session ACL and takes the next higher priority after global policy rules.

The predefined role session ACL has the following properties:

- It cannot be deleted through the WebUI or CLI. It is only deleted automatically when the corresponding role is deleted.
- It always remains at position 2 in every role and its position cannot be modified.
- It contains only application rules.
- It can be modified using the WebUI, CLI, or dashboard on a Mobility Master, however any modification results in the regeneration of ACE's for that role.
- It cannot be applied to any other role.

Each application has an implicit set of ports that are used for communication. In phase 1, if an application ACE entry is hit, the traffic matching this application's implicit port is allowed (as governed by the application ACE). The DPI engine can monitor the exchange on these ports and determine the application. Once the application is determined, phase 2 occurs when an evaluation is done to determine the final outcome for the session.

The following procedure describes how to configure the ACL application-specific parameters:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Policies** tab on the WebUI.
2. Click **+** to create a new policy. Enter a **Profile name** and select a type from the **Policy type** drop-down list and click **Submit**.
3. Select the policy created and click **+** in the **Policy<name of the rule> Rules** tab.
4. Select **Access Control** option in the **Rule Type** field.
5. Click **OK**.
6. Select IPv4 or IPv6 from the **IP version** drop-down list.
7. Select Service from the **Service/app** drop-down list and an alias from the **Service alias** drop-down list.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures the ACL application-specific parameters:

```
(host)[md](config)#ip access-list
```

Example

This example shows a DPI rule along with a L3 or L4 rule with forwarding action in the same ACL. Both ACL policies can be applied to a single user role.

ACL Policy "AppRules", Policy Type: Session

- Rule 1
  - source: any
  - destination: any
  - service or application: application facebook
  - action: permit
  - TOS value: 45
- Rule 2:
  - source: any
  - destination: any
  - service or application: application YouTube
  - action: deny

- Rule 3:
  - source: any
  - destination: any
  - service or application: application category peer-to-peer
  - action: deny
- Rule 4:
  - source: any
  - destination: any
  - service or application: TCP 23
  - action: permit
- Rule 5:
  - source: network 40.1.0.0/16
  - destination: any
  - service or application: TCP 80
  - action: permit
  - TOS: 60
- Rule 6:
  - source: network 20.1.0.0/16
  - destination: any
  - service or application: TCP 80
  - action: source-nat

ACL Policy "NetRules", Policy Type: Session

- Rule 1
  - source: network 80.0.0.0/24
  - destination: any
  - service or application: TCP 80
  - action: deny
- Rule 2:
  - source: network 60.0.0.0/24
  - destination: any
  - service or application: TCP 80
  - action: dual-nat <nat_pool>
- Rule 3:
  - source: network 10.0.0.0/24
  - destination: any
  - service or application: TCP 80
  - action: destination nat

Configuring Bandwidth Contracts

Bandwidth contract configuration lets you configure bandwidth contracts for both the global or application-specific levels.

Global Bandwidth Contract Configuration

To configure bandwidth contracts to limit application and application categories on an application or global level, or to show global bandwidth contract configuration output, execute the following commands.

```
[host] [md](config) #dpi global-bandwidth-contract[app|appcategory]
[host] [md] #show dpi global-bandwidth-contract
```

Role-Specific Bandwidth Contracts

Application-specific bandwidth contracts (unlike "generic" bandwidth-contracts) allow you to control or reserve rates for specific applications only on a per-role basis. An optional exclude list is provided that allows you to exclude applications or application categories on which a generic user or role bandwidth-contract is not applied.

Use an exclude to give specific enterprise mission-critical applications priority over other user traffic. An enterprise may have well known applications such as Microsoft Exchange, SAP, Oracle, accounting and finance applications, and other enterprise resource planning or customer relationship management applications.

Instead of enumerating bandwidth limits for each application individually on a per-user or per-role basis, you can configure a single bandwidth contract on a per-user or per-role to limit all non-mission critical applications. You can then exclude all mission-critical applications by placing them in an exclude list. This way all mission-critical applications will not be rate-limited. Important points regarding bandwidth contracts include:

- Application bandwidth contracts are per-role by default.
- When an application bandwidth-contract is configured for both a category and an application within the category, always apply the most specific bandwidth contract.

The following procedure describes how to configure role specific bandwidth contracts:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Roles** tab.
2. Click **+** to create a new user role and enter a **Name** for the new role and click **Submit** or select a role from the **Roles** table to modify an existing role.
3. Click **Show Advanced View**.
4. Select the **Bandwidth** tab.
5. Expand the **Pre-Application Limits for This Role** accordion and click + to add an application or application category to a bandwidth contract.
    a. Select the application bandwidth type from the **Type** drop-down list.
    b. Select the name of the bandwidth contract from the **Name** drop-down list.
    c. Enter values in Kbits or Mbits in the **Upstream** and **Downstream** fields.
    d. Click **Submit**.
6. Expand the **Pre-Application Limit Exceptions for This Role** accordion and click + to add an exception.

a. Select a value from the **Type** drop-down list.

b. Select an application or application category from the **Name** drop-down list.

c. Click **Submit**.



Make sure that the **Enable Deep Packet Inspection** option is checked.

7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the bandwidth application-specific parameters:

```
(host) [md] (config)# user-role <string>
(host) [md](config-role)# bw-contract exclude
```

Upgrading Application Signatures

Qosmos provides an upgraded app set library to Aruba. This is integrated in to the ArubaOS image. This is provided to the user as an Aruba-certified **proto bundle** file. The user can copy this file to Mobility Master flash and activate the **proto bundle** file using the command **dpi proto-bundle activate <filename>**.

The **proto bundle** file activation is available only from CLI.

This ensures that managed device is able to recognize the latest app set dynamically.



The proto bundle file activation must be done under managed device.

In a typical Mobility Master deployment, if the managed device is running a higher version of the proto bundle , then upgrade with a lower version will not take effect.

Protocol Database Image Upgrade

The following procedure describes the protocol database image upgrade:

1. Qosmos DPI IxEngine provides a new protocol library to Aruba.
2. Aruba uses this to create an Aruba-certified **proto bundle** file, which is provided to the user.
3. User copies this to the Mobility Master flash.
4. Under the managed device, activate the **proto bundle** file using the **dpi proto-bundle activate <filename>** command.



If you are running ArubaOS 8.0.0.0, do not upgrade the QOSMOS application set library to the latest proto bundle.

Defining Custom Application

A custom application can be created on the fly. Creating custom applications is supported on the managed device This facilitates the user to apply roles and policies, and BW contracts to the custom applications. Custom applications can be associated with custom application categories.

A maximum of 64 custom applications can be created. In each custom application, a maximum of 16 rules can be applied. A custom application can be deleted only after deleting all the rules applied on it.



Starting from ArubaOS 8.0.1.0, when a custom application is added, modified, or deleted, it takes 2 minutes for the changes to take effect.

Creating Custom Application

The following procedure describes how to create a custom application:



If you are upgrading ArubaOS to version 8.4.0.0, delete all custom applications that were created before ArubaOS 8.3.0.0. Re-create the custom applications after upgrading ArubaOS to version 8.4.0.0.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Applications** tab.
2. Expand the **Custom Application** accordion.
3. Click **+** to create a custom application.
4. Enter a **Name** for the custom application.
5. Enter an **Application ID** between 1 and 64.
6. Select a **Category** from the drop-down list, if required.
7. For **Server name**, click **+**. In the **Add Server** window:
   a. Enter the **Server name**.
   b. Enter the **URI**.
   c. Click **OK**.
8. For **Referer name**, click **+**. In the **Add Referer** window:
   a. Enter the **Referer name**.
   b. Click **OK**.
9. For **Common name**, click **+**. In the **Add Common Server** window:
   a. Enter the **Common name**.
   b. Click **OK**.
10. Click **Submit**.
11. Click **Pending Changes**.
12. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands creates a custom application:

- http host or server name based application
```
(host) [md] (config)dpi custom-app <appname> <appID>
(host) [md] (config-submode)#http <hostname>
```
- http referrer based application
```
(host) [md] (config)dpi custom-app <appname> <appID>
(host) [md] (config-submode)#http referer-param <referer>
```



Ensure that you enter only the domain name of the application for **<referer>**.

- http server name and uri based application
```
(host) [md] (config)dpi custom-app <appname> <appID>
```

```
(host) [md] (config-submode)#http <hostname-param> <hostname> uri-param <uri>
```

- https common name based application

```
(host) [md] (config)dpi custom-app <appname> <appID>
(host) [md] (config-submode)#https common-name <common-name>
```



Enter the CN of the server certificate of the application. **<app id>** is a number between 1 and 64.

The following **show** commands are introduced as part of the custom application feature:

- Issue the following commands on the Mobility Master:
  - **show dpi custom-app all**: Displays output of custom applications
  - **show dpi custom-app <appname>**: Displays the rules of custom applications.
- Issue the following commands on the managed device:
  - **show dpi custom-app all** : Displays output of custom applications.
  - **show dpi custom-app <appname>**: Displays the rules of custom applications.
  - **show dpi application custom-app all**: Displays the custom application port information and DPI application id of all the custom applications.
  - **show dpi application custom-app <appname>**: Displays the custom application port information and DPI application ID of a particular custom application.

Defining Custom Application Category

Creating user-defined custom application categories is supported on the Mobility Master. This will enable to customers to apply a policy for this category so that multiple custom applications associated with this category can receive the same policy.

A maximum of 32 custom application categories can be created.



Standard applications cannot be associated with custom application categories. Only custom applications can be associated with custom application categories. By default, custom applications fall under web category.

The following procedure describes how to create a custom application category:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Applications** tab.
2. Expand the **Custom Application** accordion.
3. Click **+** to create a custom application.
4. Select + from **Category** drop-down list.
5. In the **Application Categories** window, click **+** to create a custom application category.
6. In the **Application Categories > New Category** table:
   a. Enter a **Name** for the custom application category.
   b. Enter a **Category ID** between 1 and 32.
   c. (Optional) For **Application**, select the check box next to the list of custom applications to associate with the category. Multiple custom applications can also be selected.



The **Application** list with check box appears only if custom applications are already created.

d. Click **Submit**.

The new custom category is now available in the **Category** drop-down list.

7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command defines the application category:

```
(host) [md] (config) #dpi appcategory <appcategory> <categoryId>
```

categoryId is a number between 1 and 32.

The following CLI command associates the application category to a custom-application:

```
(host) [md] (config) #dpi custom-app <appname> <appID>
(host) [md] (config-submode)#appcategory <appcategory>
(host) [md] (config-submode)#end
```

Debugging

The following **show** commands are introduced as part of the custom application category feature:

- Issue the following command on the Mobility Master:
  - **show dpi application category user-defined all**: Displays custom app categories.
- Issue the following commands on the managed device:
  - **show dpi application category user-defined all**: Displays all custom application categories.
  - **show dpi application category user-defined <category-name>**: Displays the custom applications which associated to a particular custom application category.

Netdestination and Netservice Aliases

A netdestination is an alias for a specific host, network, or a combination of both. To use netdestination, an IP address should be configured for the host or network.

Aliases are useful for allowing or blocking specific host, network, or both. When you have multiple hosts or networks to whitelist or blacklist, you can create a single alias and add the list of hosts or network's IP addresses to it. This helps in allowing or blocking multiple entries at the same time.

Aliases can simplify configuration of session ACLs, as you can use an alias when specifying the traffic source and/or destination in multiple session ACLs. Once you configure an alias, you can use it to manage network and host destinations from a central configuration point, because all policies that reference the alias will be updated automatically when you change the alias. You can also set aliases for network protocols using netservice aliases.

The following procedure describes how to create a Netdestination alias:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Roles & Policies** > **Aliases** tab.
2. Click **+** to create a **Network Alias**.
3. Select an **IP Version** from the drop-down list.
4. Enter the **Name** for the host or domain within 63 characters.
5. Enter a **Description** of the destination within 128 characters.

6. Select **Invert** to specify that the inverse of the network addresses configured are used.
7. Click **+** to create **Items**. In the **Add New Destination Add New User Rule** window:
    a. Select a **Rule type** from the drop-down list.
    b. Enter the **IP address** if the **Rule type** is **Host**.
    c. Enter the **Domain name** if the **Rule type** is **Name**.
    d. Enter the **Start IP address** and **End IP address** if the **Rule type** is **Range**.
    e. Enter the **IP address** and **Network mask** if the **Rule type** is **Network**.
    f. Select **Vlan** from the drop-down list if the **Rule type** is **Override**.
    g. Click **OK**.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following procedure describes how to create a Netservice alias:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Roles & Policies** > **Aliases** tab.
2. Click **+** to create a **Service Alias**.
3. Enter a **Service name** for the alias within 63 characters.
4. Select a **Protocol** from the drop-down list.
    a. For TCP or UDP, select the **Port type**, **Starting port**, **End port**, and **Port list**.
5. Select **Protocol** from the **Protocol** drop-down list and enter the IP number.
6. Select an **ALG** from the drop-down list.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Table 86: *Netdestination and Netservices Parameters*

Parameter	Description
IP Version	Specifies whether the alias applies to IPv4 or IPv6 traffic. Default: IPv4
Name	Name for the host or domain. The maximum length for host name is 63 characters.
Description	Description about the destination. The maximum length of the description is 128 characters.
Invert	Specifies that the inverse of the network addresses configured are used. For example, if a network of 172.16.0.0 255.255.0.0 is configured, this parameter specifies that the alias matches everything except this subnetwork.
Rule type	Specifies the rule type applied to the alias. The rule type can be host, name, range, network, or override. Default: Override
IP address	IP address assigned to the alias.

Table 86: *Netdestination and Netservices Parameters*

Parameter	Description
Domain name	Domain name assigned to the alias name.
Start IP address	Starting IP address for a range.
End IP address	Ending IP address for a range.
Network mask	The network mask that has to be set for the alias.
VLAN	Identification number of the VLAN.
Service name	Name for the service alias.
Protocol	Configures the IP protocol value. You can configure TCP, UDP, or protocol from the drop-down list.
Port type	Select a list or range from the drop-down. Port type can be configured only for TCP or UDP.
Starting port	Sets the starting port number for a defined port range between 0 to 65535.
End port	Sets the ending port number or a defined port range between 0 to 65535.
Port list	Specifies a single port number, a list, or a defined port range by specifying both the lower and upper port numbers.
Protocol	Specify a number from 0 to 255 to define the IP protocol number.
ALG	Specify an ALG for this alias. Select one of the following service types from the drop-down list: ■ ftp: Service is FTP ■ tftp: Service is TFTP ■ dns: Service is DNS ■ dhcp: Service is DHCP ■ sip: Service is SIP ■ sips: Service is Secure SIP ■ svp: Service is SVP ■ sccp: Service is SCCP ■ rtsp: Service is RTSP ■ vocera: Service is VOCERA ■ noe: Service is Alcatel NOE ■ h323: Service is H323 ■ jabber: Service is JABBER ■ facetime: Service is facetime

IP Classification-based Firewall

ArubaOS supports IP classification-based firewall. IP classification helps to identify the IP address and geolocation from where malicious activities originate.

With the IP classification any inbound attack from the malicious end points may be stopped at the managed device itself and thereby, protect the client devices behind the managed device. IP

classification uses the IP reputation and IP geolocation databases. The IP reputation and IP geolocation databases are periodically updated and synchronized with partnering servers from Webroot or Brightcloud.

The IP reputation database contains all the current known IP addresses associated with various malicious activities. This database lists the IP addresses and the corresponding threats, like botnet, DoS, spam sources, and so on originated by them. If an IP address is classified as malicious, the traffic sent to or received from that address may be denied based on the configured policy.

The IP geolocation database determines the geographical location of an IP address from where the traffic is received or to which the traffic is sent. If the geographic location of an IP address is determined, traffic may be permitted or denied after scanning the configured geography-based rules policy.

The IP geolocation database provides granularity of the geolocation of the IP address to the level of a country and city. It provides powerful visualizations that may be used to demonstrate the top countries that originate malware or spyware traffic or top countries from where maximum DDoS attacks are received. This information may be used to formulate geolocation firewall policies to protect the internal network resources and keep the network healthy.

The benefits of this IP classification-based firewall include:

- Identify and prevent any attack from a malicious host.
- Identify the geolocation of incoming or outgoing traffic.
- Identify the location from where maximum spyware, malware, or DDoS attacks originate.
- Provide geolocation visibility information about the traffic flows.
- Formulate firewall policies based on geolocation of IP address to permit or deny traffic.

The existing firewall policy enforcement in ArubaOS relies on L3 or L4 to L7 information with DPI or WebCC support. IP classifiaction extends the firewall by allowing a user to define new IP classification-based firewall policies.



- All traffic originating from Remote AP users is exempted from location-based firewall policies.
- IP classification is applied to all traffic destined to the managed device for all forwarding modes, except the bridge mode traffic which is locally routed at the AP. For split-tunnel mode, IP classification is applied only for the split-tunneled traffic that is destined to the managed device.
- IP classification is applicable only for IPv4 addresses. IP classification based access policies for IPv6 addresses is not supported.

To enable IP classification-based firewall using the CLI:

```
(host) [mynode] (config) #firewall
(host) [mynode] (config-submode)#ip-classification
```

To configure IP classification-based policy using the CLI:

```
(host) [mynode] (config) #ip access-list geolocation global-geolocation-acl [permit|deny]
[to|from] location
```

where location is either of:

- anonymous_proxy - Match packets from or to anonymous proxy
- any - Match any location
- country - Match packets from or to a country
- region - Match packets from or to a region

To configure IP reputation rule using the CLI:

```
(host) [mynode] (config) #ip-reputation deny [inbound|outbound]
```

To view IP reputation table using the CLI:

```
(host) [mynode] #show datapath ip-reputation
```

To view IP reputation counters using the CLI:

```
(host) [mynode] #show datapath ip-reputation counters
```

To view IP reputation real time cache using the CLI:

```
(host) [mynode] #show datapath ip-reputation rtc
```

To view IP geolocation table using the CLI:

```
(host) [mynode] #show datapath ip-geolocation
```

To view IP geolocation counters using the CLI:

```
(host) [mynode] #show datapath ip-geolocation counters
```

To view IP classification table using the CLI:

```
(host) [mynode] #show datapath session ip-classification
```

To view IP classification-based policy using the CLI:

```
(host) [mynode] #show ip access-list global-geolocation-acl
```

Chapter 23
ClearPass Policy Manager Integration

ArubaOS and ClearPass Policy Manager include support for centralized policy definition and distribution. ArubaOS now supports downloadable user roles for both wired and wireless users in cluster deployments. By using this feature, when ClearPass Policy Manager successfully authenticates a user, the user is assigned a role by ClearPass Policy Manager. If the role is not defined on the Mobility Master, the role attributes can also be downloaded automatically.

ArubaOS now provides IPv6 support for ClearPass Policy Manager downloadable user role that allows you to configure ACL and policy enforcement profile in the ClearPass Policy Manager. The managed device can download the user role using IPv6 address configured in RADIUS authentication server. Hence, the managed device can now use either IPv4 or IPv6 address to download the user role from ClearPass Policy Manager.



The IPv6 address support is applicable from CPPM 6.9.0 or later versions.

This chapter contains the following sections:

- Introduction on page 529
- Important Points to Remember on page 529
- Enabling Downloadable Role on a Managed Device on page 530
- Sample Configuration on page 531
- Per-Command Authorization for Management Users on page 535

Introduction

In order to provide highly granular per-user level access, user roles can be created when a user has been successfully authenticated. During the configuration of a policy enforcement profile in ClearPass Policy Manager, the administrator can define a role that should be assigned to the user after successful authentication. In RADIUS authentication, when ClearPass Policy Manager successfully authenticates a user, the user is assigned a role by ClearPass Policy Manager. If the role is not defined on Mobility Master, the role attributes can also be downloaded automatically. This feature supports roles obtained by the following authentication methods:

- 802.1X (WLAN and wired users)
- MAC authentication
- Captive Portal

Important Points to Remember

- Under **Advanced** mode, ClearPass Policy Manager does not perform any error checking to confirm accuracy of the role definition. Therefore, it is recommended that you review the role defined in ClearPass Policy Manager prior to enabling this feature.

- The attributes that are listed below, herein referred to as whitelist role attributes, can be defined in ClearPass Policy Manager:
  - **netdestination**
  - **netservice**
  - **ip access-list eth**
  - **ip access-list mac**
  - **ip access-list session**
  - **user-role**
- The above attributes that are referred to by a role definition must either be defined within the role definition itself or configured on the Mobility Master before the policy is downloaded.
- In ClearPass Policy Manager, two or more attributes (as listed above) should not have the same name. The following example is considered invalid, as both the attributes use **test** as the profile or net destination name:
  ```
  qos-profile test
  netdestination test
  ```
- Instance names (name of a whitelist role attribute) are case-sensitive. Attributes must adhere to the following rules:
  - Should not match any CLI option nested under a command from the whitelist.
  - Should not contain a number or a combination of numbers.
  - Should not contain any periods '.'.
  - Should not contain any spaces.

The example below is considered an invalid configuration and prevents ClearPass Policy Manager role download on a managed device:

```
netservice 'tcp' tcp 443
```

The first instance of **tcp** is a user-defined field, while the second is an operator of the **netservice** command. This violates the first rule.

```
netdestination 'alias'
```

The user-defined name **alias** is also a valid operator of the **netdestination** command. This violates the first rule.

```
netdestination '10.1.5'
```

This user-defined name uses both numbers and periods. This violates the second and third rule.

```
ip access-list stateless '100'
```

This user-defined name uses numbers. This violates the second rule.

```
qos-profile emp role
```

This profile name **emp role** contains spaces. This violates the fourth rule.



It is recommended that some naming convention similar to the CamelCase (mixture of upper and lower case letters in a single word) be used to avoid collisions with the CLI options in the role description.

Enabling Downloadable Role on a Managed Device

The following procedure describes how to enable downloadable role on a managed device:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication** page.

2. Select the **AAA Profiles** tab
3. Expand **AAA** in the **AAA Profiles** list, and then select a **AAA** profile.
4. Select the **Download Role from CPPM** check box to enable role download.
5. Click **Submit**.
6. Select **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands enable downloadable role on a managed device:

```
(host) [md] (config) #aaa profile <profile-name>
(host) [md] (AAA profile) #download-role
```

Sample Configuration

The following example shows the configuration details to integrate ClearPass Policy Manager server with a managed device to automatically download roles.

ClearPass Policy Manager Server Configuration

This section describes the following topics:

Adding a Device

The following procedure describes how to add a device:

1. Navigate to the **Configuration > Network > Devices** page in the ClearPass Policy Manager server.
2. Click **Add** above the **Network Devices** list. The **Add Device** page opens.
3. Under the **Device** tab, enter the **Name**, **IP or Subnet Address**, and **RADIUS Shared Secret** fields.
4. Keep the rest of the fields as default.
5. Click **Add**.

The fields are described in Table 87.

Table 87: *Device Tab*

Parameter	Description
Name	The name or identity of the device.
IP or Subnet Address	The IP address or subnet (example 10.1.1.1/24) of the device.
RADIUS Shared Secret	Enter and confirm a Shared Secret for each of the two supported request protocols.

Adding an Enforcement Profile

The following procedure describes how to add an enforcement profile:

1. Navigate to the **Configuration > Enforcement > Profiles** page.
2. Click **Add** above the **Enforcement Profiles** list. The **Enforcement Profiles** page opens.
3. Under the **Profile** tab, select **Aruba Downloadable Role Enforcement** from the **Template** drop-down list.

4. Enter the **Name** of the enforcement profile.
5. Under **Role Configuration Mode**, select **Advanced**.
   Keep the rest of the fields as default.
6. Click **Next**.

For the rest of the configuration, see Advanced Role Configuration Mode.

The fields are described in Table 88.

Table 88: *Enforcement Profiles Page*

Parameter	Description
Template	Policy Manager comes pre-packaged with several enforcement profile templates. In this example, select **Aruba Downloadable Role Enforcement** - RADIUS template that can be filled with user role definition to create roles that can be assigned to users after successful authentication.
Name	The name of the enforcement profile.
Role Configuration Mode	▪ **Standard**: Configures the enforcement profile role using standard mode. ▪ **Advanced**: Configures the enforcement profile role using advanced mode.

Advanced Role Configuration Mode

The following procedure describes how to enable advanced role configuration mode:

1. Under the **Attributes** tab, select **Radius:Aruba** from the **Type** table.
2. From the **Name** drop-down list, select **Aruba-CPPM-Role**.
3. In the **Value** field, enter the attribute for the downloadable-role.
4. Click the **Save** icon to save the attribute.
5. Click **Save** to save the enforcement profile.

The fields are described in Table 89.

Table 89: *Enforcement Profiles Attributes Tab*

Parameter	Description
Type	Any RADIUS vendor dictionary that is pre-packaged with Policy Manager, or imported by the Administrator. This field is pre-populated with the dictionary names.
Name	The name of the attribute from the dictionary selected in the Type field. The attribute names are pre-populated from the dictionary.
Value	The attribute for the downloadable role. You can enter free-form text to define the role and policy. **NOTE:** The maximum limit for free form text is 16,000 bytes.

Adding Enforcement Policy

The following procedure describes how to add an enforcement policy:

1. Navigate to the **Configuration > Enforcement > Policies** page.
2. Click **Add** above the **Enforcement Policies** list. The **Enforcement Policies** page opens.
3. Under the **Enforcement** tab, enter the **Name** of the enforcement policy.
4. From the **Default Profile** drop-down list, select **[Deny Access Profile]**.

   Keep the rest of the fields as default.
5. Click **Next**.

The fields are described in Table 90.

Table 90: *Enforcement Policies Enforcement Tab*

Parameter	Description
Name	The name of the enforcement policy.
Default Profile	An Enforcement Policy applies Conditions (roles, health, and time attributes) against specific values associated with those attributes to determine the Enforcement Profile. If none of the rules matches, Policy Manager applies the Default Profile. See Adding an Enforcement Profile on page 531 to add a new profile.

6. Under the **Rules** tab, click **Add Rule**. The **Rules Editor** page opens.
7. Select **Click to add...** from the **Conditions** section. Select the appropriate values, and then click the **Save** icon.
8. In the **Enforcement Profiles** section, select the RADIUS enforcement profile that you created in Adding an Enforcement Profile on page 531from the **Profile Names** drop-down list.
9. Click **Save**.

The fields are described in Table 91.

Table 91: *Enforcement Policies Rules Editor*

Parameter	Description
Type	The rules editor appears throughout the Policy Manager interface. It exposes different namespace dictionaries depending on Service type. When working with service rules, you can select **Authentication** namespace dictionary
Name	Drop-down list of attributes present in the selected namespace. In this example, select **Source**.
Operator	Drop-down list of context-appropriate (with respect to the attribute) operators. In this example, select **EQUALS**.
Value	Drop-down list of the Authentication source database. In this example, select **[Local User Repository]**.
Profile Names	Name of the RADIUS enforcement profile.

Adding Services

The following procedure describes how to add services:

1. Navigate to the **Configuration > Services** page.
2. Click **Add** above the **Services** list.

3. Under the **Service** tab, select **802.1X Wired** from the **Type** drop-down-list.
4. In the **Name** field, enter the name of the service.

   Keep the rest of the fields as default.
5. Click **Next**.

The fields are described in Table 92.

Table 92: *Service Tab*

Parameter	Description
Type	The service type. In this example, select **802.1X Wired**.
Name	The name of the service.

6. Under the **Authentication** tab, select **[Local User Repository] [Local SQL DB]** from the **Authentication Sources** drop-down list.

   Keep the rest of the fields as default.
7. Click **Next** twice.
8. Under the **Enforcement** tab, select the enforcement policy that you created in Adding Enforcement Policy on page 532 from the **Enforcement Policy** drop-down list.

   Keep the rest of the fields as default.
9. Click **Save**.

For more configuration details on ClearPass Policy Manager, see the *ClearPass Policy Manager User Guide.*

Managed Device Configuration

This section describes the following topics:

Configuring ClearPass Policy Manager Server on a Managed Device

The following CLI commands configure ClearPass Policy Manager server on a managed device:

```
(host) [md] (config) #aaa authentication-server radius cppm_server
(host) [md] (RADIUS Server "cppm_server") #host <ip_address_of_
  cppm_server>
(host) [md] (RADIUS Server "cppm_server") #key <psk>
(host) [md] (RADIUS Server "cppm_server") #cppm username <username>
  password <password>
```

Configuring Server Group to include ClearPass Policy Manager Server

The following CLI commands configure a server group to include ClearPass Policy Manager server:

```
(host) [md] (config) #aaa server-group cppm_grp
(host) [md] (server group "cppm_grp") #auth-server cppm_server
```

Configuring 802.1X Profile

The following CLI command configures a 802.1X profile:

```
(host) [md] (config) #aaa authentication dot1x cppm_dot1x_prof
```

Configuring AAA Profile

The following CLI commands configure a AAA profile:

```
(host) [md] (config) #aaa profile cppm_aaa_prof
```

```
(host) [md] (AAA Profile "cppm_aaa_prof") #authentication-dot1x cppm_
  dot1x_prof
(host) [md] (AAA Profile "cppm_aaa_prof") #dot1x-server-group cppm_gr
 (AAA Profile "cppm_aaa_prof") #download-role
```

Show AAA Profile

The following CLI command displays an AAA profile:

```
(host) [md] #show aaa profile cppm_aaa_prof
  AAA Profile "cppm_aaa_prof"
  ---------------------
  Parameter                          Value     Set
  ---------                          -----     ---
  Initial role                       logon
  MAC Authentication Profile         N/A
  MAC Authentication Default Role    guest
  MAC Authentication Server Group    default
  802.1X Authentication Profile      N/A
  802.1X Authentication Default Role guest
  802.1X Authentication Server Group N/A
  Download Role from CPPM            Disabled
  Set username from dhcp option 12   Disabled
  L2 Authentication Fail Through     Disabled
  Multiple Server Accounting         Disabled
  User idle timeout                  N/A
  Max IPv4 for wireless user         2
  RADIUS Accounting Server Group     N/A
  RADIUS Interim Accounting          Disabled
  XML API server                     N/A
  RFC 3576 server                    N/A
  User derivation rules              N/A
  Wired to Wireless Roaming          Enabled
  Device Type Classification         Enabled
  Enforce DHCP                       Disabled
  PAN Firewall Integration           Disabled
  Open SSID radius accounting        Disabled
```

For additional command parameters, see the *ArubaOS CLI Reference Guide*.

Per-Command Authorization for Management Users

Starting from ArubaOS 8.6.0.0, ArubaOS supports per-command authorization for management users with TACACS+ Servers running on CPPM. This feature gives flexibility in determining commands to be allowed for each management user at each configuration-node. The allowed and not-allowed commands for each management user can be configured in the TACACS+ servers. The commands executed by the management user (with a certain administrative role) will be sent to the TACACS+ server for authorization and only the authorized commands can be executed. Otherwise, the command triggered will be denied.

For TACACS+ Server running on a CPPM, a new management role **"tacacs-authz"** needs be chosen as the Aruba-Admin-Role in **"Aruba:Common"** service-type for the target TACACS Enforcement Profile for the authenticated management-user.

For more information, please refer to the TACACS+ Enforcement Profiles section in the ClearPass Policy Manager 6.8 User Guide.



This feature is available only on TACACS+ servers running on ClearPass Policy Manager.

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 BSSID which is usually the AP's MAC address.

In the Aruba network, an AP uses a unique BSSID for each WLAN, so each individual AP or AP group can support multiple WLAN configurations.

Basic WLAN Configuration

The recommended method for creating a new WLAN configuration is through the new WLAN wizard, although advanced users may also configure a WLAN manually.

Creating a WLAN using the WLAN Wizard

To start the New WLAN wizard, in the **Managed Network** node hierarchy, navigate to **Configuration > Tasks** and select **Create a new WLAN**. The wizard opens and prompts you to enter the following information:

Configuration Setting	Description
General	
Name (SSID)	Name you assign to the new WLAN.
Primary usage	Select whether the WLAN will be primarily supporting employees or guest users.
Broadcast on	Choose whether the WLAN SSID should broadcast on all APs associated to the managed device or Mobility Master configuration, or whether the WLAN should broadcast on APs in a selected AP group. If you choose the **Select AP Groups** option, you are prompted to select one or more AP groups.
Forwarding mode	If the forwarding mode is set to **Tunnel**, data is tunneled to the managed device using GRE. When a WLAN is configured to use the **Decrypt-Tunnel** forwarding mode, that AP decrypts and decapsulates all 802.11 frames from a client and sends the 802.3 frames through the GRE tunnel to the managed device, which then applies firewall policies to the user traffic. When the managed device sends traffic to a client, the managed device sends 802.3 traffic through the GRE tunnel to the AP, which then converts it to encrypted 802.11 and forwards to the client.
VLANs	
VLAN	The VLAN(s) into which users are placed in order to obtain an IP address. If you are creating a guest WLAN, remember that guest users must be separated from employee users by VLANs in the network.

Configuration Setting	Description
Named VLANs	Click **Show VLAN Details** to view the list of named VLANs configured on the managed device or Mobility Master. To add a new VLAN, click **+** in the **Named VLANs** table, then enter appropriate values in the following fields: **VLAN name:** Name for the new VLAN **VLAN ID/Range:** Specify the beginning and ending VLAN IDs separated by a hyphen. For example, 55-58. To edit a named VLAN, select the VLAN from the table and click the pencil button. You can edit the **VLAN name** and **VLAN ID/Range** parameters.
VLAN IDs	Select a **VLAN** from the **Named VLANs** table to view the list of VLAN IDs configured on the managed device or Mobility Master. To add a new VLAN ID, click **+** in the **VLAN IDs** table, then enter/select appropriate values in the following fields: **VLAN ID:** Identification number for the VLAN **Admin state:** Enable or disable the VLAN interface. To edit a VLAN ID, select the VLAN from the **VLAN IDs** table and click the pencil button. You can edit the **VLAN ID** and **Admin state** settings.
Security*(for employee WLANs)*	
Enterprise	This option supports the following configuration parameters: **Key management**: Use this setting to select the layer-2 encryption type to be used on this WLAN SSID. Select either **WPA-3 Enterprise** (default), **WPA-2 Enterprise**, or **WPA Enterprise**. **Use CNSA suite**: Use Commercial National Security Algorithm (CNSA) for enterprise network. **Auth servers**: To add an existing server, Click **+** to open the **Add Existing Server** window and select a preconfigured server from the list of servers. To define a new server, click + on the Add Existing Server window and define a new **LDAP** or **RADIUS** server. For details, see Configuring Authentication Servers on page 193 **Reauth interval**: Define interval, in seconds or minutes, between re-authentication attempts in either minutes or seconds. **MAC authentication**: Select this option to enforce machine authentication before user authentication. If selected, either the machine-default-role or the user-default-role is assigned to the user, depending on which authentication is successful. **Blacklisting**: Blacklists the client if authentication fails a specified number of times. **Max authentication failures**: If **Blacklisting** is enabled, this parameter defines the number of times a user can try to login with wrong credentials after which the user is blacklisted as a security threat.
Personal	This option supports the following configuration parameters: **Key management**: Use this setting to select the layer-2 encryption type to be used on this WLAN SSID. Select either **WPA-3 Personal** (default), **WPA-2 Personal**, or **WPA Personal**. **Passphrase**: Enter a password for the WLAN. **Retype:** Retype the password. **MAC authentication**: Select this option to enforce machine authentication before user authentication. If selected, either the machine-default-role or the user-default-role is assigned to the user, depending on which authentication is successful. **Blacklisting**: Blacklists the client if authentication fails a specified number of times.

Configuration Setting	Description
	Max authentication failures: If **Blacklisting** is enabled, this parameter defines the number of times a user can try to login with wrong credentials after which the user is blacklisted as a security threat.
Open	This option supports the following configuration parameters: **MAC authentication**: Select this option to enforce machine authentication before user authentication. If selected, either the machine-default-role or the user-default-role is assigned to the user, depending on which authentication is successful.
Security*(for guest WLANs)*	
ClearPass or other external captive portal	This option supports the following configuration parameters: **Auth servers**: Click **+** to open the **Add Existing Server** window and select a preconfigured server from the list of servers. To define a new server, click + in the **Add Existing Server** window and define a new **LDAP** or **RADIUS** server. For details, see Configuring Authentication Servers on page 193 **CPPM host**: IPv4 address of the ClearPass Policy Manager host. **CPPM page**: URL of the page that appears for the user logon. This can be set to any URL. Default: /auth/index.html. **Redirect URL**: URL to which an authenticated user will be directed. This parameter must be an absolute URL that begins with either http:// or https://.
Internal captive portal with authentication	This option supports the following configuration parameters: **Template**: Define the title, text, banner icon and banner color for the captive portal landing page. **Redirect URL**: URL to which an authenticated user will be directed. This parameter must be an absolute URL that begins with either http:// or https://. **Custom HTML**: Click this link to browse to and select HTML files for the initial login and welcome pages.
Internal captive portal with email registration	This option supports the following configuration parameters: **Template**: Define the title, text, banner icon and banner color for the captive portal landing page. **Redirect URL**: URL to which an authenticated user will be directed. This parameter must be an absolute URL that begins with either http:// or https://. **Custom HTML**: Click this link to browse to and select HTML files for the initial login and welcome pages.
Internal captive portal, no auth or registration	This option supports the following configuration parameters: **Template**: Define the title, text, banner icon and banner color for the captive portal landing page. **Redirect URL**: URL to which an authenticated user will be directed. This parameter must be an absolute URL that begins with either http:// or https://. **Custom HTML**: Click this link to browse to and select HTML files for the initial login and welcome pages.
No Captive Portal	Guests are granted access without a captive portal.
Access	
Default role	Select a user role to be assigned to an employee that successfully authenticates to the WLAN.

Configuration Setting	Description
	If you are creating an employee WLAN, click the **Default role** drop-down list and select an existing user role, or define a new role for the WLAN, by clicking on **Show Roles** and clicking **+** in the **Roles** table. If you are creating a guest WLAN, the WLAN wizard automatically creates a default role for the guest users that have successfully authenticated to the WLAN, named is **<WLAN-name>-guest-logon**. To configure this role, in the **Managed Network** node hierarchy, navigate to **Configuration > Roles & Policies > Roles** tab and select the created role. As you configure your guest role, keep in mind the following guidelines for guest WLANs: Guests must be limited not only in where they may go, but also by what network protocols and ports they may use to access resources. Guests should be allowed to access only the local resources that are required for IP connectivity. These resources include DHCP and possibly DNS if an outside DNS server is not available. In most cases, a public DNS is always available. All other internal resources should be off limits for the guest. This restriction is achieved usually by denying any internal address space to the guest user. A time-of-day restriction policy should be used to allow guests to access the network only during normal working hours, because they should be using the network only while conducting official business. A rate limit can also be put on each guest user to keep the user from using up the limited wireless bandwidth. Accounts should be set to expire when their local work is completed, typically at the end of each business day. For complete information on creating user roles and assigning rules and policies to a role, see Roles and Policies on page 491
Server-derived roles	(For employee WLANs using enterprise security) Enable this option to configure server derivation rules that can assign a 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.
Derivation method	(For employee WLANs using enterprise security) Select a derivation method. Select **Use value returned from ClearPass or other auth server** if your users will authenticate to the WLAN via ClearPass Policy Manager or another type of authentication server, or select **User rules defined in table below** to define a custom role based upon RADIUS Server VSAs. Click **+** in the **Role Derivations Rules** table and define the following values: **Attribute**: RADIUS VSA type **Condition**: contains, equals, not-equals, start-with or value-of **Operand**: Text string compared against VSA condition **Role**: Role assigned if the VSA condition and operand match. For the current and complete list of all RADIUS VSAs available in the version of ArubaOS currently running on your managed device, access the command-line interface and issue the command **show aaa radius attributes**. See also Configuring Authentication Servers on page 193

The following procedure describes how to manually configure a WLAN that uses 802.1X authentication.



This method for configuring a WLAN is recommended for advanced users only.

1. Configure your authentication servers.
2. Create an authentication server group, and assign the authentication servers you configured in step 1 to that server group.

3. Configure a firewall access policy for a group of users
4. Create a user role, and assign the firewall access policy you created in step 3 to that user role.
5. Configure the AAA profile for the configuration node.
    a. Assign the user role defined in step 4 to the **802.1X Authentication Default Role** associated to the AAA profile.
    b. Associate the server group you created in step 2 to the AAA profile.
6. Configure the SSID profile for the configuration node.
7. Configure the virtual AP profile for the configuration node, the Virtual AP profile for the configuration node will automatically be associated to the AAA profile configured in Step 5, and the SSID profile configured in Step 6.

The following CLI commands configure a WLAN:



This method for configuring a WLAN is recommended for advanced users only.

```
(host)[node](config) #aaa server-group "THR-DOT1X-SERVER-GROUP-WPA2"
auth-server Internal
!
ip access-list session THR-POLICY-NAME-WPA2
user any any permit
!
(host)[node](config) #user-role THR-ROLE-NAME-WPA2
session-acl THR-POLICY-NAME-WPA2
!
(host)[node](config) #aaa server-group "THR-DOT1X-SERVER-GROUP-WPA2"
auth-server Internal
!
(host)[node](config) #aaa profile "THR-AAA-PROFILE-WPA2"
dot1x-default-role "THR-ROLE-NAME-WPA2"
dot1x-server-group "THR-DOT1X-SERVER-GROUP-WPA2"
!
(host)[node](config) #wlan ssid-profile "THR-SSID-PROFILE-WPA2"
essid "THR-WPA2"
opmode wpa2-aes
!
(host)[node](config) #wlan virtual-ap "THR-VIRTUAL-AP-PROFILE-WPA2"
ssid-profile "THR-SSID-PROFILE-WPA2"
aaa-profile "THR-AAA-PROFILE-WPA2"
vlan 60
!
(host)[node](config) #ap-group "THRHQ1-STANDARD"
virtual-ap "THR-VIRTUAL-AP-PROFILE-WPA2"
```

WLAN Configuration Profiles

You can configure your WLANS 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.

When you define a WLAN using the New WLAN wizard on the **Configuration > Tasks** page of the Mobility Master or stand-alone controller WebUI, the wizard automatically creates a new virtual AP profile, AAA profile, 802.1X, Server group profile and SSID profile with the same name as the WLAN, and

with the configuration settings and values defined via the wizard. These profiles also support additional advanced features that are not configurable via the WLAN wizard on the **Configuration > Tasks** page.

The following table describes the profiles that comprise the configuration settings for an ArubaOS WLAN, with links to the sections of this document that describe these profiles in more detail.

Table 93: *WLAN Profiles*

Profile	Description
Virtual AP Profile	This is the top-level WLAN configuration profile. A Virtual AP profile allows you to configure WLAN settings such as broadcast/multicast settings, forwarding modes and RF bands, but it also identifies the individual 802.11k, AAA, Anyspot, Hotspot 2.0, SSID and WWM Traffic management profiles to be used by that WLAN. Default profile name: <WLAN Name> When you create a WLAN using the WLAN wizard, ArubaOS automatically creates a new Virtual AP profile with the same name as the WLAN.
802.11k profile	The 802.11k protocol provides mechanisms for APs and clients to dynamically measure the available radio resources. Each 802.11k profile also references one instance of each the following additional profile types. ■ Beacon Report Request profile: Defines beacon report request settings. Beacon report requests are sent only to 802.11k-compliant clients that advertise Beacon Report Capability in their Radio Resource Management Enabled Capabilities IE. ■ Radio Resource Management IE profile: Defines Radio Resource Management Information Elements for WLANs with 802.11k support enabled. ■ Traffic Stream Measurement Report Request profile: Defines Traffic Stream Measurement report requests. These report requests are sent only to 802.11k- compliant clients that advertise a traffic stream report capability. Default profile name: default
AAA profile	The AAA profile defines the type of authentication used by clients associating to a WLAN. Each AAA profile also references one instance of each the following additional profile types: ■ 802.1X Authentication profile: Defines 802.1X authentication settings. ■ 802.1X Authentication Server Group profile: Defines fail through and load balancing settings for a group of servers used for 802.1X authentication. ■ MAC Authentication profile: Defines MAC authentication settings. ■ MAC Authentication Server Group profile: Defines fail through and load balancing settings for a group of servers used for MAC authentication. ■ RADIUS Accounting Server Group profile: Defines fail through and load balancing settings for a group of servers used for RADIUS accounting. ■ RFC 3576 Server profile: Defines a RADIUS server to send user disconnect, CoA, and session timeout messages as described in RFC 3576. ■ XML API Server profile: Define an authentication key for an XML API server, to perform customized external captive portal user management using an XML API interface. Default profile name: <WLAN Name> When you create a WLAN using the WLAN wizard, ArubaOS automatically creates a new AAA profile with the same name as the WLAN.
AnySpot Profile	The Anyspot client probe suppression feature decreases network traffic by suppressing probe requests from clients attempting to locate and connect to other known networks. By default, a virtual AP is not associated with an Anyspot profile, so an Anyspot profile must first be defined, and then manually associated to the virtual AP. Default profile name: N/A

Profile	Description
Hotspot 2.0 Profile	Hotspot 2.0 is a WFA Passpoint specification based upon the 802.11u protocol that provides wireless clients with a streamlined mechanism to discover and authenticate to suitable networks, and allows mobile users the ability to roam between partner networks without additional authentication. The Hotspot profile for a WLAN references a hotspot *advertisement* profile, which in turn references several other profiles that define settings for individual hotspot features. ■ Hotspot Advertisement Profiles ■ ANQP Venue Name Profiles ■ ANQP Network Authentication Profiles ■ ANQP Domain Name Profile ■ ANQP IP Address Availability Profiles ■ ANQP NAI Realm Profiles ■ ANQP Roaming Consortium Profiles ■ ANQP 3GPP Cellular Network Profiles ■ H2QP Connection Capability Profiles ■ H2QP Operator Friendly Name Profiles ■ H2QP Operating Class Indication Profiles ■ H2QP WAN Metrics Profiles Default profile name: <WLAN Name> When you create a WLAN using the WLAN wizard, ArubaOS automatically creates a new Hotspot 2.0 profile with the same name as the WLAN.
SSID Profile	A SSID profile defines the name of the network, authentication type for the network, basic rates, transmit rates, SSID cloaking, and certain WMM settings for the network. Each SSID profile also references one instance of each the following additional profile types: ■ 802.11r profile: The Fast BSS Transition (802.11r) mechanism minimizes the delay when a voice client transitions from one BSS to another within the same ESS. ■ EDCA Parameters (AP) profile: ArubaOS supports media access prioritization through EDCA, which defines four ACs to prioritize traffic. This profile defines EDCA settings for APs. ■ EDCA Parameters (Station) profile: ArubaOS supports media access prioritization through EDCA, which defines four ACs to prioritize traffic. This profile defines EDCA settings for clients. ■ High-throughput SSID profile: Defines 802.11ac very-high-throughput settings for the 5 GHz frequency band, and high-throughput (802.11n) settings for both the 5 GHz and 2.4 GHz frequency bands. ■ High-efficiency SSID profile: Defines 802.11ax spectrum efficiency and area throughput on both the 2.4 GHz and 5 GHz frequency bands. Default profile name: <WLAN Name> When you create a WLAN using the WLAN wizard, ArubaOS automatically creates a new SSID profile with the same name as the WLAN.

Starting from ArubaOS 8.0.1.0, you can modify the parameters of profiles that are associated to a WLAN.



The Virtual AP profile parameters cannot be modified.

Modifying Profile Parameters Associated with WLANs

The following procedure describes how to modify profile parameters associated with WLANs:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > WLANs** page.
2. Select a WLAN in the **WLANs** table and click on the **Profiles** tab.
3. In the **Profiles for WLAN <WLAN name>** list, select **Wireless LAN > Virtual AP > <WLAN name>**.
4. Select an associated profile from the list and make the necessary changes.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Configuring the Virtual AP Profile

The recommended method for creating a new WLAN configuration is through the new WLAN wizard, although advanced users may also configure a WLAN manually.



For important information on changing the virtual AP forwarding mode for a WLAN serving active wired or wireless clients, see Changing a Virtual AP Forwarding Mode on page 550.

Manually Configuring the Virtual AP Profile

The following procedure describes how to configure Virtual AP profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN > Virtual AP**.
3. To edit an existing Virtual AP profile, select the Virtual AP profile you want to edit. To create a new Virtual AP profile, click **+** and enter a name for the new Virtual AP profile in the **Profile name** field.
   The Virtual AP profile settings are divided into four sections: **General, RF, Advanced,** and **Broadcast/Multicast.**
4. Configure your Virtual AP settings, the profile parameters in each section are described in Virtual AP Profile Parameters.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Table 94: *Virtual AP Profile Parameters*

Parameter	Description
General	
Virtual AP enable	Select the **Virtual AP enable** check box to enable or disable the virtual AP.
VLAN	The VLAN(s) into which users are placed in order to obtain an IP address. **NOTE:** You must add an existing VLAN ID to the Virtual AP profile.

Parameter	Description
Forward mode	This parameter controls whether data is tunneled to the managed device using GRE, bridged into the local Ethernet LAN (for remote APs), or a combination thereof depending on the destination (corporate traffic goes to the managed device, and Internet access remains local). All forwarding modes support band steering, TSPEC/TCLAS enforcement, 802.11k and station blacklisting. Click the drop-down list to select one of the following forward modes: ▪ **Tunnel:** The AP handles all 802.11 association requests and responses, but sends all 802.11 data packets, action frames and EAPOL frames over a GRE tunnel to the managed device for processing. The managed device 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 managed device) 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 managed device before you configure campus APs in bridge mode. **NOTE:** In a bridge mode, the wired or wireless clients which have the same IP address as the Remote AP's local DHCP server cannot communicate with other devices even if the AP is deployed as a Campus AP. If you want to use the default Remote AP's IP address as the client IP address, you need to change the Remote AP's DHCP server IP address to a different IP address. To change Remote AP's DHCP server IP address, see Enabling Remote AP Advanced Configuration Options on page 859. ▪ **Split-Tunnel:** 802.11 frames are either tunneled or bridged, depending on the destination (corporate traffic goes to the managed device, and Internet access remains local). A Remote AP in split-tunnel forwarding mode handles all 802.11 association requests and responses, encryption/decryption, and firewall enforcement. the 802.11e and 802.11k action frames are also processed by the remote AP, which then sends out responses as needed. ▪ **Decrypt-Tunnel:** Both remote and campus APs can be configured in decrypt-tunnel mode. When an AP uses decrypt-tunnel forwarding mode, that AP decrypts and decapsulates all 802.11 frames from a client and sends the 802.3 frames through the GRE tunnel to the managed device, which then applies firewall policies to the user traffic. When the managed device sends traffic to a client, the managed device 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 managed device.

Parameter	Description
	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 are not present for APs in regular tunnel forwarding mode. You must enable the control plane security feature on the managed device before you configure campus APs in decrypt-tunnel forward mode. **NOTE:** Virtual APs in bridge or split-tunnel mode using static WEP should use key slots 2-4 on the managed device. Key slot 1 should only be used with Virtual APs in tunnel mode.
RF	
Allowed band	The band(s) on which to use the virtual AP: ■ **g**—802.11b/g band only (2.4 Ghz). ■ **a**—802.11a band only (5 Ghz). ■ **all**—both 802.11a and 802.11b/g bands (5 GHz and 2.4 GHz). This is the default setting.
Band Steering	ARM's band steering feature encourages dual-band capable clients to stay on the 5 GHz band on dual-band APs. This frees up resources on the 2.4 Ghz 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 5 Ghz band than on the 2.4 Ghz band. Dual-band 802.11n-capable clients may see even greater bandwidth improvements, because the band steering feature will automatically select between 40MHz or 20 Mhz channels in 802.11n networks. This feature is disabled by default, and must be enabled in a Virtual AP profile. The band steering feature supports both campus APs and remote APs that have a virtual AP profile set to tunnel, split-tunnel or bridge forwarding mode. Note, however, that if a campus or remote APs has virtual AP profiles configured in bridge or split-tunnel forwarding mode but no virtual AP in tunnel mode, those APs will gather information about 5G-capable clients independently and will not exchange this information with other APs that also have bridge or split-tunnel virtual APs only.
Steering Mode	Band steering supports the following three different band steering modes. ■ **Balance-bands:** In this band steering mode, the AP tries to balance the clients across the two radios in order to best utilize the available 2.4G bandwidth. This feature takes into account the fact that the 5 Ghz band has more channels than the 2.4 Ghz band, and that the 5 Ghz channels operate in 40 Mhz while the 2.5 Ghz band operates in 20 MHz. ■ **Prefer-5GHz** (default): If you configure the AP to use **prefer-5GHz** band steering mode, the AP will try to steer the client to 5G band (if the client is 5G capable) but will let the client connect on the 2.4G band if the client persists in 2.4 G association attempts. ■ **Force-5GHz:** When the AP is configured in **force-5GHz** band steering mode, the AP will try to force 5Ghz-capable APs to use that radio band.

Parameter	Description
Advanced	
Cellular handoff assist	When both the ClientMatch and the cellular handoff assist features are enabled, the cellular handoff assist feature can help a dual-mode, 3G or 4G-capable Wi-Fi device such as an iPhone, iPad or Android client at the end of a Wi-Fi network switch from Wi-Fi to an alternate 3G or 4G radio that provides better network access. This feature is supported by iOS and Android devices only.
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.
Blacklist Time	Number of seconds that a client is quarantined from the network after being blacklisted. Default: 3600 seconds (1 hour)
Deny inter user traffic	Select this check box to deny traffic between the clients using this virtual AP profile. The global firewall shown in the **Configuration>Advanced Services > Stateful Firewall > Global** window also includes an option to deny all inter-user traffic, regardless of the Virtual AP profile used by those clients. If the global setting to deny inter-user traffic is enabled, all inter-user traffic between clients will be denied, regardless of the settings configured in the virtual AP profiles. If the setting to deny inter-user traffic is disabled globally but enabled on an individual virtual AP, only the traffic between un-trusted users and the clients on that particular virtual AP will be blocked. **NOTE:** This field is not applicable across controllers even when they are in the same cluster.
Deny time range	Click the drop-down list and select a configured time range for which the AP will deny access. If you have not yet configured a time range, navigate to **Configuration > Security > Access Control > Time Ranges** to define a time range before configuring this setting in the Virtual AP profile.
DoS Prevention	If enabled, APs ignore de-authentication frames from clients. This prevents a successful de-authorization attack from being carried out against the AP. This does not affect third-party APs. Default: Disabled
HA Discovery on-association	If enabled, home agent discovery is triggered on client association instead of home agent discovery based on traffic from client. Mobility on association can speed up roaming and improve connectivity for clients that do not send many uplink packets to trigger mobility (VoIP clients). Best practices is to disable this parameter as it increases IP mobility control traffic between managed devices in the same mobility domain. Enable this parameter only when voice issues are observed in VoIP clients. Default: Disabled **NOTE: ha-disc-onassoc** parameter works only when IP mobility is enabled and configured on the managed device. For more information about this parameter, see Home Agent Discovery on Association on page 814
Mobile IP	Enables or disables IP mobility for this virtual AP. Default: Enabled

Parameter	Description
Preserve Client VLAN	If you select this check box, clients retain their previous VLAN assignment if the client disassociates from an AP and then immediately re-associates either with same AP or another AP on the same managed device.
Remote-AP Operation	Configures when the virtual AP operates on a remote AP: ■ **standard** (default)—Enables the virtual AP when the remote AP connects to the managed device. This option can be used for any (bridge/split-tunnel/tunnel/d-tunnel) virtual APs. ■ **persistent**—Permanently enables the virtual AP after the remote AP initially connects to the managed device (Bridge Mode only). This option can be used for any (Open/PSK/802.1X) bridge VAPs. ■ **backup**—Enables the virtual AP if the remote AP cannot connect to the managed device (Bridge Mode only). This option can be used for non-802.1X bridge VAPs. ■ **always**—Permanently enables the virtual AP (Bridge Mode only). This option can be used for non-802.1X bridge VAPs.
Station Blacklisting	Select this check box to enable detection of DoS attacks, such as ping or SYN floods, that are not spoofed deauthorization attacks. Default: Enabled
Strict Compliance	If enabled, the AP denies client association requests if the AP and client station have no common rates defined. Some legacy client stations which are not fully 802.11-compliant may not include their configured rates in their association requests. Such non-compliant stations may have difficulty associating with APs unless strict compliance is disabled. Default: Disabled
VLAN Mobility	Enable or disable VLAN (Layer-2) mobility. Default: Disabled
WAN operation mode	This feature works in conjunction with the WAN Health Check Manager and Uplink Manager. When all uplinks are down, the uplink manager makes the needed changes based on configuration and pushes these changes to APs. ■ If the operation mode is set to **primary**, the VAP will be disabled. ■ If the operation mode is set to **backup**, the VAP will be enabled. ■ If the operation mode is set to **always**, the VAP will not change.
FDB Update on Assoc	This parameter enables seamless failover for silent clients, allowing them to re-associate. If you select this option, the controller will generate a Layer 2 update on behalf of client to update forwarding tables in bridge devices. Default: Disabled
Broadcast/Multicast	
Dynamic Multicast Optimization (DMO)	Enable/Disable dynamic multicast optimization. This parameter is disabled by default, and cannot be enabled without the PEFNG license.
Dynamic Multicast Optimization (DMO) Threshold	Maximum number of high-throughput stations in a multicast group beyond which dynamic multicast optimization stops. Range: 2-255 stations

Parameter	Description
	Default: 6 stations.
Drop Broadcast and Multicast	Select the **Drop Broadcast and Multicast** check box to filter out broadcast and multicast traffic in the air. Do not enable this option for virtual APs configured in bridge forwarding mode. This configuration parameter is only intended for use for virtual APs in tunnel mode. In tunnel mode, all packets travel to the controller, so the controller is able to drop all broadcast traffic. When a virtual AP is configured to use bridge forwarding mode, most data traffic stays local to the AP, and the controller is not able to filter out that broadcast traffic. **IMPORTANT:** If you enable this option, you must also enable the **Convert Broadcast ARP requests to unicast** parameter on the virtual AP profile to prevent ARP requests from being dropped.
Convert Broadcast ARP requests to unicast	If enabled, all broadcast ARP requests are converted to unicast and sent directly to the client. You can check the status of this option using the **show ap active** and the **show datapath tunnel** command. If enabled, the output will display the letter **a** in the flags column. This configuration parameter is only intended for use for virtual APs in tunnel mode. In tunnel mode, all packets travel to the controller, so the controller is able to convert ARP requests directed to the broadcast address into unicast. When a virtual AP is configured to use bridge forwarding mode, most data traffic stays local to the AP, and the controller is not able to convert that broadcast traffic. This parameter is enabled by default. Behaviors associated with these settings are enabled upon upgrade to ArubaOS 6.1.3.2. If your controller supports clients behind a wireless bridge or virtual clients on VMware devices, you must disable this setting to allow those clients to obtain an IP address. In previous releases of ArubaOS, the virtual AP profile included two unique broadcast filter parameters; the **drop broadcast and multicast** parameter, which filtered out all broadcast and multicast traffic in the air except DHCP response frames (these were converted to unicast frames and sent to the corresponding client) and the **convert Broadcast ARP requests to unicast** parameter, which converted broadcast ARP requests to unicast messages sent directly to the client. The **Convert Broadcast ARP requests to unicast** setting includes the additional functionality of broadcast-filter all parameter, where DHCP response frames are sent as unicast to the corresponding client. This can impact DHCP discover/requested packets for clients behind a wireless bridge and virtual clients on VMware devices. Disable this option to resolve this issue and allow clients behind a wireless bridge or VMware devices to receive an IP address. Default: Enabled

A Virtual AP profile directly references one of each of the following profiles types:

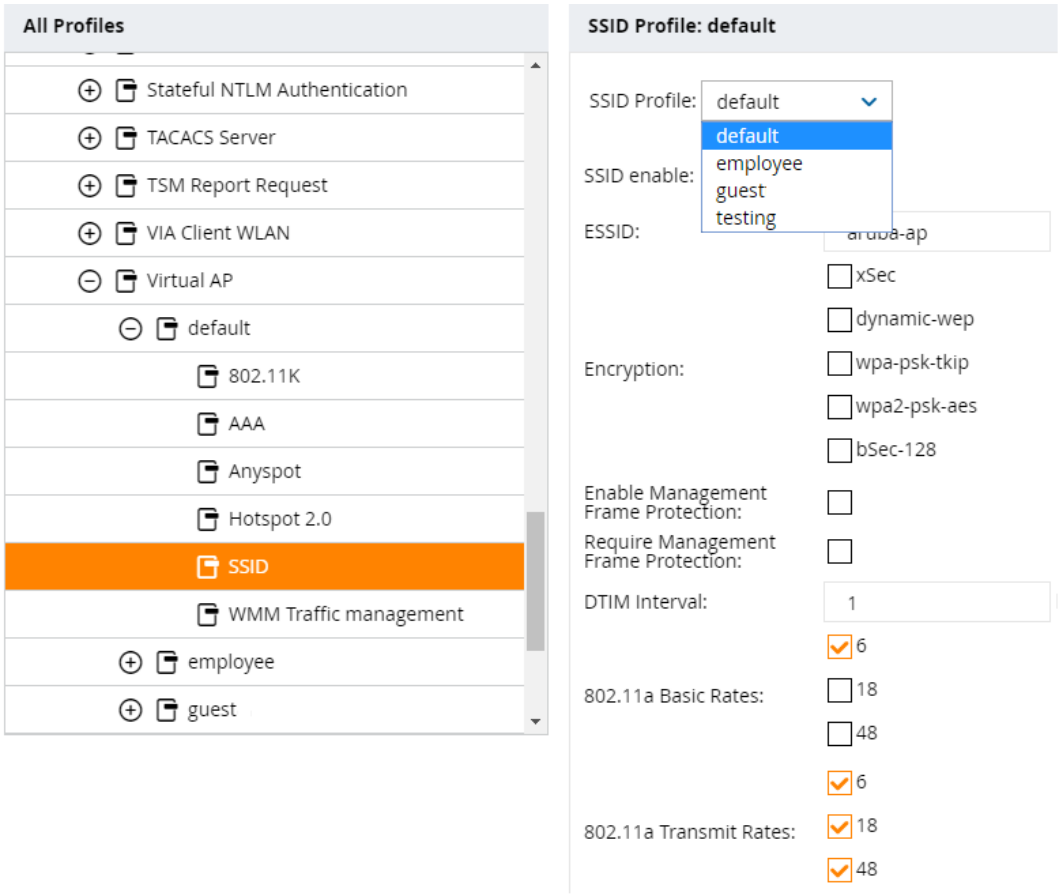
- 802.11k
- AAA
- AnySpot
- HotSpot 2.0
- SSID
- WWM Traffic Management

The following procedure describes how to change the profiles associated to a Virtual AP profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN > Virtual AP**.

3. Select the Virtual AP profile you want to edit. The **All Profiles** window displays the list of associated profiles for that Virtual AP.
4. Select any of the associated profiles in the list.
5. A drop-down list appears at the top of the right window pane which allows you to select another profile for that type.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Figure 84 *Associating Profiles to a Virtual AP*



The following CLI commands configure a Virtual AP profile:

```
(host)[node](config) #wlan virtual-ap <profile>
(host)[node] (Virtual AP profile "profile")aaa-profile <profile>
(host)[node] (Virtual AP profile "profile")anyspot-profile <profile>
(host)[node] (Virtual AP profile "profile")dot11k-profile <profile>
(host)[node] (Virtual AP profile "profile")hs2-profile <profile>
(host)[node] (Virtual AP profile "profile")ssid-profile <profile>
(host)[node] (Virtual AP profile "profile")wmm-traffic-management-profile <profile>
```

Modifying Profiles and Parameters Associated with AP Groups

The following procedure describes how to modify profiles and parameters associated with AP groups:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > AP Groups** page.
2. Select an AP group in the **AP Groups** table and click on the **Profiles** tab.

3. Select a profile under **Profiles for Group <AP Group>**.
4. Click **<NAME> profile** drop-down list and select a profile.
5. Make the necessary changes to the profile and click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Selective Multicast Streams

The selective multicast group is based only on the packets learned through the Internet Group Management Protocol (IGMP).

- When the **Drop Broadcast and Multicast** setting is enabled in the virtual AP profile, the managed device allows multicast packets to be forwarded only if the following conditions are met:
  - packets originating from the wired side have a destination address range of 225.0.0.0 - 239.255.255.255
  - a station has subscribed to a multicast group.
- If the **DMO** setting is enabled in the virtual AP profile , the packets are sent with 802.11 unicast header.
- When IGMP snooping/proxy is disabled, the managed device is not aware of the IGMP membership and drops the multicast flow.
- If AirGroup is enabled, mDNS (SSDP) packets are sent to the AirGroup application. The common address for mDNS is 224.0.0.251 and SSDP is 239.255.255.250.

Changing a Virtual AP Forwarding Mode

When you change the forwarding mode for a Virtual AP actively serving clients, the user table will NOT reflect accurate client information unless the entries for those users are manually cleared. Use the following procedure to change the forwarding mode on a Virtual AP serving wired or wireless clients.

Changing the Forwarding Mode for Wired Users

To change the forwarding mode for wired users connected to the wired port on an AP:

1. Disable the port by issuing the CLI command **ap wired-port-profile <ap-wired-port-profile>** shutdown.

   This will disconnect any wired clients using that port.
2. Issue the command **aaa user delete {<ipaddr>|all|mac <macaddr>|name <username>|role <role>}** to remove from the user table the wired users associated with AP wired ports using the <ap-wired-port-profile>.
3. Issue the command **ap wired-ap-profile <profile> forward-mode <mode>** where **<mode>** is the new forwarding mode for the wired port.
4. Re-enable the port using the command **ap wired-port-profile <ap-wired-port-profile> no shutdown**.

Changing the Forwarding Mode for Wireless Users

To change the forwarding mode for wireless users associated with an AP radio:

1. Issue the command **ap-name <group> no virtual-ap <vap-profile>** or **ap-group <group> no virtual-ap <vap-profile>** to disassociate the AP or group of APs from the virtual AP profile.

2. Issue the command **aaa user delete {<ipaddr>|all|mac <macaddr>|name <username>|role <role>}** to remove from the user table the users associated to the virtual-ap specified in the previous step.
3. Issue the command **wlan virtual-AP <vap-profile> forward-mode <mode>** where **<mode>** is the new forwarding mode for the virtual AP.
4. Issue the command **ap-name <group> virtual-ap <vap-profile>** or **ap-group <group> virtual-ap <vap-profile>** to reassociate the AP or group of APs with the virtual AP profile.

Radio Resource (802.11k) and BSS Transition Management (802.11v)

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 802.11v BSS Transition capability can improve throughput, data rates and QoS for the voice clients in a network by shifting (via transition) the individual voice traffic loads to more appropriate points of association within the ESS.

Configuring the 802.11k Profile

The following procedure describes how to configure the 802.11k profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN>802.11k**.
3. To edit an existing 802.11k profile, select the 802.1k profile you want to edit. To create a new 802.11k profile, click **+** and enter a name for the new 802.11k profile in the **Profile name** field.
4. Configure your 802.11k radio settings. The configuration parameters are described in Table 95.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Table 95: *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. Enabling this option also enables support for the 802.11v BSS transition management feature described in BSS Transition Management (802.11v) on page 558. Default: Disabled
Forcefully disassociate on-hook voice clients	Select this option to allow the AP to forcefully disassociate *on-hook* voice clients (clients that are not on a call) after period of inactivity. Without the forced disassociation feature, if an AP has reached its call admission control limits and an on-hook voice client wants to start a new call, that client may be denied. If forced disassociation is enabled, those clients can associate to a neighboring AP that can fulfill their QoS requirements. Default: Disabled

Parameter	Description
Measurement Mode for Beacon Reports	Select any one of the following measurement modes from the drop down list: ■ **active-all-ch**—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. ■ **active-ch-rpt**—In this mode, the client and returns a report that contains a list of channels in a regulatory class where a client is likely to find an AP, including the AP transmitting the AP channel report. ■ **beacon-table** (default)—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 does not support the selected measurement mode, it returns a Beacon Measurement Report with the incapable bit set in the Measurement Report Mode field.
Channel for Beacon Requests in 'A' band	This value is sent in the 'Channel' field of the beacon requests on the 'A' radio. You can specify values in the range 34 to 165. Default: 36
Channel for Beacon Requests in 'BG' band	This value is sent in the 'Channel' field of the Beacon Requests on the 'BG' radio. You can specify values in the range 1 to 14. Default: 1
Channel for AP Channel Reports in 'A' band	This value is sent in the 'Channel' field of the AP channel reports on the 'A' radio. You can specify values in the range 34 to 165. Default: 36.
Channel for AP Channel Reports in 'BG' band	This value is sent in the 'Channel' field of the AP channel reports on the 'BG' radio. You can specify values in the range 1 to 14. Default: 1
Time duration between consecutive Beacon Requests	This option configures the time duration between two consecutive beacon requests sent to a 802.11k client. However, if a different value is required, the `bcn-req-time` option can be used. This permits values in the range from 10 seconds to 200 seconds. A value of 0 is used to indicate that the generation of Beacon Request frames is turned off. Default: 60
Time duration between consecutive Link Measurement Requests	This option configures the time duration between two consecutive link measurement requests sent to a 802.11k client. This parameter permits values in the range from 10 seconds to 200 seconds. A value of 0 is used to indicate that the generation of Link Measurement Request frames is turned off. Default: 60

Parameter	Description
Time duration between consecutive Transmit Stream Measurement Requests	This option configures the time duration between two consecutive transmit stream measurement requests sent to a 802.11k client. This permits values in the range from 10 seconds to 200 seconds. A value of 0 is used to indicate that the generation of Transmit Stream Measurement Request frames is turned off. Default: 90

The following CLI command configures configure 802.11k profiles. The available parameters for this profile are described in Table 95.

```
(host)[node](config)#wlan dotllk-profile <profile-name>
```

Configuring Radio Resource Management Information Elements

ArubaOS supports the following radio resource management (RRM) information elements for APs with 802.11k support enabled.

The following procedure describes how to select the radio resource management information elements to be sent in beacons and probe responses:

1. Navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN > RRM IE.**
3. To edit an existing RRM IE profile, select the RRM IE profile you want to edit. To create a new RRM IE profile, click **+** and enter a name for the new RRM IE profile in the **Profile name** field.
4. Configure your RRM IE settings. The configuration parameters are described in Table 96.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Table 96: *RRM IE Parameters*

Parameter	Description
Advertise Enabled Capabilities IE	This value is used to determine if the RRM Enabled Capabilities IE should be advertised in the beacon frames. A value of "Enabled" allows the RRM Enabled Capabilities IE to be present in the beacon frames when 802.11K capability is enabled. A value of "Disabled" prevents the advertisement of the RRM Enabled Capabilities IE in the beacon frames when 802.11K capability is enabled.
Advertise Country IE	This value is used to determine if the Country IE should be advertised in the beacon frames. A value of "Enabled" allows the Country IE to be present in the beacon frames when 802.11K capability is enabled. A value of "Disabled" prevents the advertisement of the Country IE in the beacon frames when 802.11K capability is enabled.
Advertise Power Constraint IE	This value is used to determine if the Power Constraint IE should be advertised in the beacon frames. A value of "Enabled" allows the Power Constraint IE to be present in the beacon frames when 802.11K capability is enabled. A value of "Disabled" prevents the advertisement of the Power Constraint IE in the beacon frames when 802.11K capability is enabled.

Parameter	Description
Advertise TPC Report IE	This value is used to determine if the TPC Report IE should be advertised in the beacon frames. A value of "Enabled" allows the TPC Report IE to be present in the beacon frames when 802.11K capability is enabled. A value of "Disabled" prevents the advertisement of the TPC Report IE in the beacon frames when 802.11K capability is enabled.
Advertise QBSS Load IE	This value is used to determine if the QBSS Load IE should be advertised in the beacon frames. A value of "Enabled" allows the QBSS Load IE to be present in the beacon frames when 802.11K capability is enabled. A value of "Disabled" prevents the advertisement of the QBSS Load IE in the beacon frames when 802.11K capability is enabled.
Advertise BSS AAC IE	This value is used to determine if the BSS Available Admission Capacity IE should be advertised in the beacon frames. A value of "Enabled" allows the BSS Available Admission Capacity IE to be present in the beacon frames when 802.11K capability is enabled. A value of "Disabled" prevents the advertisement of the BSS Available Admission Capacity IE in the beacon frames when 802.11K capability is enabled.
Advertise Quiet IE	This value is used to determine if the Quiet IE should be advertised in the beacon frames. A value of "Enabled" allows the Quiet IE to be present in the beacon frames when 802.11K capability is enabled. A value of "Disabled" prevents the advertisement of the Quiet IE in the beacon frames when 802.11K capability is enabled.

The following CLI command configures radio resource management information elements in the RRM IE profile:

(config)#wlan rrm-ie-profile <profile>

Configuring Beacon Report Requests

The beacon report requests are sent only to 802.11k-compliant clients that advertise Beacon Report Capability in their RRM Enabled Capabilities IE. The beacon request frames are sent every 60 seconds.

The following procedure describes how to select the information to be sent in beacon report requests:

1. Navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN > Beacon Report Requests.**
3. To edit an existing Beacon Report Request profile, select the Beacon Report Request profile you want to edit. To create a new Beacon Report Request profile, click **+** and enter a name for the new Beacon Report Request profile in the **Profile name** field.
4. Configure your Beacon Report Request settings.

   The configuration parameters are described in Table 97.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Table 97: *Beacon Report Request Parameters*

Parameter	Description
Interface	This field is used to specify the Radio interface for transmitting the Beacon Report Request frame. It can have a value of either 0 or 1. Default: 1.
Regulatory Class	This option is used to specify the Regulatory Class field in the Beacon Report Request frame. It can be set to one of the following: - ■ **1** (for 5 GHz band) ■ **12** (for 2.4 GHz band)
Channel	This option is used to set the Channel field in the Beacon Report Request frame. The Channel value can be set to one of the following: - the channel of the AP (when Measurement Mode is set to either 'Passive' or 'Active-All channels') - 0 (when Measurement Mode is set to 'Beacon Table') - 255 (when Measurement Mode is set to 'Active-Channel Report')
Randomization Interval	This value is used to set the Randomization Interval field in the Beacon Report Request frame. The Randomization Interval is used to specify the desired maximum random delay in the measurement start time. It is expressed in units of TUs (Time Units). A Randomization Interval of 0 in a measurement request indicates that no random delay is to be used. This field can be given a value in the range (0, 65535). Default: 0
Measurement Duration	This value is used to set the Measurement Duration field in the Beacon Report Request frame. The Measurement Duration is set to the duration of the requested measurement. It is expressed in units of Time Units. This field can be given a value in the range (0, 65535). Default: 0
Measurement Mode for Beacon Reports	Select one of the following measurement modes from the drop down list: ■ **active-all-ch-**—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. ■ **active-ch-rpt**—In this mode, the client and returns a report that contains a list of channels in a regulatory class where a client is likely to find an AP, including the AP transmitting the AP channel report. ■ **beacon-table** (default)—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

Parameter	Description
	end of the measurement duration, compiles all received beacons or probe response with the requested SSID and BSSID into a measurement report. **NOTE:** If a station doesn't support the selected measurement mode, it returns a Beacon Measurement Report with the Incapable bit set in the Measurement Report Mode field. Default Mode: beacon-table
Reporting Condition	This option is used to indicate the value for the "Reporting Condition" field in the Beacon Reporting Information sub-element present in the Beacon Report Request frame. It can have a range from 0 to 255. Default: 0
ESSID name	This option is used to indicate the value for the "SSID" field in the Beacon Report Request frame. It corresponds to the SSID Name for which the Beacon Report Request frame needs to be generated. It is a string with a minimum length of 1 and a maximum length of 32.
Reporting Detail	This option is used to indicate the value for the "Detail" field in the Reporting Detail sub-element present in the Beacon Report Request frame. Default: Disabled
Measurement Duration Mandatory	This value is used to set the "Duration Mandatory" bit of the Measurement Request Mode field of the Beacon Report Request frame. Default: Disabled
Request Information values	This option is used to indicate the contents of the Request Information IE that could be present in the Beacon Report Request frame. The Request Information IE is present for all Measurement Modes except the 'Beacon Table' mode. It consists of a list of Element IDs that should be included by the client in the response frame.

The following CLI command selects the information to be sent in beacon report requests:

```
(config)wlan bcn-rpt-req-profile <profile>
```

Configuring Traffic Stream Measurement Report Requests

The Traffic Stream Measurement (TSM) report requests are sent only to 802.11k- compliant clients that advertise a traffic stream report capability. The TSM report request frames are sent every 60 seconds.

The following procedure describes how to select the information to be sent in TSM report requests.

1. Navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN > TSM Report Request.**
3. To edit an existing TSM Report Request profile, select the TSM Report Request profile you want to edit. To create a new TSM Report Request profile, click **+** and enter a name for the new TSM Report Request profile in the **Profile name** field.
4. Configure your TSM Report Request settings.

   The configuration parameters are described in Table 98.
5. Click **Submit**.

6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Table 98: *TSM Report Request Parameters*

Parameter	Description
Request Mode for TSM Report Request	Select one of the following request modes: ▪ normal ▪ triggered This value is used to determine the request mode for the Transmit Stream/Category Measurement Request frame. A Transmit Stream/Category Measurement Request frame can be sent in either normal mode or triggered mode. There are two options for this parameter **normal** and **triggered**. When the **triggered** option is selected, the Transmit Stream/Category Measurement Request frame is sent only when the trigger condition occurs. The default value for this field is **normal**.
Number of repetitions	This value is used to set the "Number of Repetitions" field in the Transmit Stream/Category Measurement Request frame. The **Number of Repetitions** field contains the requested number of repetitions for all the Measurement Request elements in this frame. A value of zero in this field indicates Measurement Request elements are executed once without repetition. A value of 65535 in the Number of Repetitions field indicates Measurement Request elements are repeated until the measurement is cancelled or superseded. This field has values in the range (0, 65535). Default: 65535.
Duration Mandatory	This value is used to set the "Duration Mandatory" bit of the Measurement Request Mode field of the Transmit Stream/Category Measurement Request frame. Default: Enabled
Randomization Interval	This value is used to set the Randomization Interval field in the Transmit Stream/Category Measurement Request frame. The Randomization Interval is used to specify the desired maximum random delay in the measurement start time. It is expressed in units of Time Units. When the request mode for the Transmit Stream/Category Measurement Request frame is set to "triggered", the Randomization Interval is not used and is set to 0. A Randomization Interval of 0 in a measurement request indicates that no random delay is to be used. This field can be given a value in the range (0, 65535). Default: 0.
Measurement Duration	This value is used to set the Measurement Duration field in the Transmit Stream/Category Measurement Request frame. The Measurement Duration is set to the duration of the requested measurement. It is expressed in units of TUs. When the request mode for the Transmit Stream/Category Measurement Request frame is set to **triggered**, the Measurement Duration field should be set to 0. This field can be given a value in the range (0, 65535). Default: 9776
Traffic ID	The value is used to set the Traffic Identifier field in the Transmit Stream/Category Measurement Request frame. The Traffic Identifier field contains the TID subfield. The TID subfield indicates the TC or TS for which traffic is to be measured. This field can be given a value in the range (0, 255). Default: 96

Parameter	Description
Bin 0 Range	This value is used to set the 'Bin 0 Range' field in the Transmit Stream/Category Measurement Request frame. Bin 0 Range indicates the delay range of the first bin (Bin 0) of the Transmit Delay Histogram, expressed in units of TUs. This field can be given a value in the range (0, 255). Default: 6

The following CLI command selects the information to be sent in TSM report requests:

(config) #wlan tsm-req-profile <profile>

BSS Transition Management (802.11v)

BSS Transition Management enables an AP to request a voice client to transition to a specific AP, or suggest a set of preferred APs to a voice client, due to network load balancing or BSS termination. This helps the voice client identify the best AP to which that client should transition to as that client roams. ArubaOS supports BSS Transition Management features defined by the 802.11v standard.



Both the 802.11v BSS transition management features and the 802.11k radio resource management features are disabled by default. To enable both of these features, select the **Advertise 802.11k Capability** option in an 802.11k profile.

Frame Types

BSS Transition Management uses the following frame types:

- **Query:** A Query frame is sent by the voice client that supports BSS transition management requesting a BSS transition candidate list to its associated AP, if the associated AP indicates that it supports the BSS transition capability.
- **Request:** An AP that supports BSS Transition Management responds to a BSS Transition Management Query frame with a BSS Transition Management Request frame. The AP may also send an unsolicited BSS Transition Management Request frame to a voice client at any time, if the client supports the BSS Transition Management capability. The Request frame also contains a Disassociation flag. If the flag is set, then the AP forcefully disassociates the client after 10 beacon intervals.
- **Response:** A Response frame is sent by the voice client back to the AP, informing whether it accepts or denies the transition.

802.11k and 802.11v clients

For 802.11k capable clients, the client management framework uses the actual beacon report generated by the client in response to a beacon report request sent by the AP. This beacon report replaces the virtual beacon report for that client. For 802.11v capable clients, the controller uses the 802.11v BSS Transition message to steer clients to the desired AP upon receiving a client steer trigger from the AP.

Enabling 802.11v BSS Transition Management

802.11v BSS transition management is enabled by default. Client match uses this feature to steer devices. The **cm-dot11v** parameter in the **rf arm-profile** enables or disables client match to use the 802.11v feature.

```
(host) ^[mynode] (config) #rf arm-profile default
(host) ^[mynode] (Adaptive Radio Management (ARM) profile "default") # [no] cm-dot11v
```

Fast BSS Transition (802.11r)

ArubaOS provides support for Fast BSS Transition as part of the 802.11r implementation. Fast BSS Transition mechanism minimizes the delay when a voice client transitions from one BSS to another within the same ESS. Fast BSS Transition establishes security and QoS states at the target AP before or during a re-association. This minimizes the time required to resume data connectivity when a BSS transition happens.

Starting from ArubaOS 8.7.0.0, the WPA3 opmodes support fast BSS transition.

The following list provides the supported opmodes:

- WPA3-Personal (SAE)
- WPA3-Personal (SAE) + Transition Enable
- WPA3-Enterprise Basic (WPA3-AES-CCM-128)
- WPA3-Enterprise Non-CNSA (WPA3-AES-GCM-256)

The following table provides the modes in which Fast BSS Transition is supported:

Table 99: *Supported VAP Forwarding Modes*

VAP Forwarding Mode	Support for 802.11r
Tunnel Mode	Yes
Decrypt-Tunnel Mode	Yes
Split-Tunnel Mode	No
Bridge Mode	Beta quality

Important Points to Remember

Fast BSS Transition is operational only if the wireless client has support for 802.11r standard. If the client does not have support for 802.11r standard, it falls back to normal WPA2 authentication method.

Configuring Fast BSS Transition

To enable and configure Fast BSS Transition on a configuration node, you must create and configure an 802.11r profile.



Fast BSS transition is operational only with WPA2-Enterprise, WPA2-Personal, WPA3-Personal, WPA3-Enterprise Basic, and WPA3-Enterprise Non-CNSA mode with GCM-256 encryption.

The following procedure describes how to configure fast BSS transition:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN** > **802.11r.**
3. To edit an existing 802.11r profile, select the 802.1r profile you want to edit. To create a new 802.11r profile, click **+** and enter a name for the new 802.11r profile in the **Profile name** field.
4. Configure your 802.11r radio settings:
    a. Select the **Advertise 802.11r Capability** option to allow Virtual APs using this profile to advertise 802.11r capability.

b. Enter the mobility domain ID value (1-65535) in the **802.11r Mobility Domain ID** field. The default value is 1.

c. Enter the R1 Key timeout value in seconds (60-86400) for decrypt-tunnel or bridge mode in the **802.11r R1 Key Duration** field. The default value is 3600.

5. Click **Submit**.



Assign the edited 802.11r profile or the new 802.11r profile to an SSID profile, otherwise the 802.11r capability cannot be used.

6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands create an 802.11r profile:

```
(host) [node] (config) #wlan dot11r-profile dot11r_profile
(host) ^[node] (802.11r Profile "dot11r_profile") #dot11r
```

Assign the 802.11r profile to an SSID profile using the following command:

```
(host) [node] (config) #wlan ssid-profile ssid_profile
(host) ^[node] (SSID Profile "ssid_profile") #dot11r-profile dot11r_profile
```

Troubleshooting Fast BSS Transition

ArubaOS provides various troubleshooting options to verify the Fast BSS Transition functionalities.

In decrypt-tunnel mode and bridge mode, each r0 key generates up to four r1 keys and the managed device pushes each r1 key to the corresponding AP. The following commands help verifying the pushing functionality:

Execute the following command to view all the r1 keys that are stored in an AP:

```
(host)[node](config) #show ap debug dot11r state
```

You can execute the following command to remove an r1 key from an AP when the AP does not have a cached r1 key during Fast BSS Transition roaming:

```
(host)[node] #ap debug dot11r remove-key
```

Execute the following command to view the hit/miss rate of r1 keys cached on an AP before a Fast BSS Transition roaming. This counter helps to verify if enough r1 keys are pushed to the neighboring APs:

```
(host)(config) #show ap debug dot11r efficiency <client-mac>
```

WLAN SSID Profiles

An SSID is the network or WLAN that any client sees. A SSID profile defines the name of the network, authentication type for the network, basic rates, transmit rates, SSID cloaking, and certain WMM settings for the network.

SSID Profile Overview

ArubaOS supports different types of the AES, TKIP, and WEP encryption. AES is the most secure and recommended encryption method. Most modern devices are AES capable and AES should be the default

encryption method. Use TKIP only when the network includes devices that do not support AES. In these situations, use a separate SSID for devices that are only capable of TKIP.

Suite-B Cryptography

The Suite-B (bSec) protocol is a pre-standard protocol that has been proposed to the IEEE 802.11 committee as an alternative to 802.11i. The main difference between bSec and standard 802.11i is that bSec implements Suite-B algorithms wherever possible. Notably, AES-CCM is replaced by AES-GCM, and the Key Derivation Function (KDF) of 802.11i is upgraded to support SHA-256 and SHA-384. In order to provide interoperability with standard Wi-Fi software drivers, bSec is implemented as a shim layer between standard 802.11 Wi-Fi and a Layer 3 protocol such as IP. A managed device configured to advertise a bSec SSID will advertise an open network, however only bSec frames will be permitted on the network.



This feature requires the ACR license.

The bSec protocol requires that you use VIA 2.1.1 or greater on the client device. Consult VIA documentation for more information on configuring and installing VIA.

The bSec protocol is available in 128-bit mode and 256-bit mode. The number of bits specifies the length of the AES-GCM encryption key. Using United States Department of Defense classification terminology, bSec-128 is suitable for protection of information up to the SECRET level, while bSec-256 is suitable for protection of information up to the TOP SECRET level.

Suite-B AES-128-GCM and AES-256-GCM encryption is supported by the ArubaOS hardware.

Wi-Fi Multimedia Protection

Wi-Fi Multimedia™ (WMM®) is a Wi-Fi Alliance® certification program that is based on the IEEE 802.11e amendment. WMM ensures QoS for latency-sensitive traffic in the air. WMM divides the traffic into four queues or access categories:

- voice
- video
- best effort
- background

Management Frame Protection

ArubaOS supports the IEEE 802.11w standard, also known as Management Frame Protection (MFP). Management Frame Protection makes it difficult for an attacker to deny service by spoofing Deauth and Disassoc management frames. Management Frame Protection uses 802.11i (Robust Security Network) framework that establishes encryption keys between the client and AP.

Management Frame Protection is configured on a virtual AP as part of the wlan ssid-profile. SSIDs that support WPA2 opmode support MFP in all forwarding mode except tunnel mode. SSIDs that support WPA3 opmode support MFP in tunnel mode only. Two MFP related parameters, **mfp-capable** and **mfp-required**, cannot be configured through the CLI or WebUI. ArubaOS automatically configures these parameters based on the opmode.



Management Frame Protection can only be enabled on SSIDs that support WPA2 or WPA3.

High-Efficiency WLAN (HE)

ArubaOS 8.4.0.0 supports the IEEE 802.11ax standard, also known as High-Efficiency WLAN (HE). HE improves spectrum efficiency and area throughput in dense deployment scenarios of APs or stations in both indoor and outdoor environments. HE enhances the 802.11 PHY and MAC channels on both 2.4 GHz and 5 GHz frequency bands.

HE includes the following features:

- Backward compatible with 802.11a/b/g/n/ac.
- Better power management for longer battery life.

HE is configured on a virtual AP as part of the WLAN SSID profile. You can configure the **High-efficiency SSID** profile from the WebUI. For details, see High-Efficiency (HE) APs on page 694High-Efficiency (HE) APs on page 694.

Multi Band Operation (MBO)

MBO enables the network to utilize the available spectrum efficiently, and helps in optimizing connectivity experience for the end-users. MBO, also known as Agile Multiband is a pre-requisite for the 802.11ax certification, therefore any AP or STA that supports 802.11ax will have the MBO capabilities.

MBO helps the APs and STAs exchange information to allow the network utilize the available spectrum efficiently. MBO works to facilitate efficient use of multiple frequency bands or channels that are available in the APs and the STAs. Starting from ArubaOS 8.6.0.0, 510 Series, 530 Series, and 550 Series access points support the Agile Multiband.

MBO can be configured using the wlan ssid profile using the WebUI. For details, see Configuring the SSID Profile on page 562

WLAN Ageout Refresh Direction

The refresh direction of an SSID profile is bidirectional by default. Starting from ArubaOS 8.5.0.0, the refresh direction of the SSID profile can be configured to use either bidirectional, receive-only, or transmit-only data frames. Bidirectional indicates data frames from both directions, receive-only indicates data frames that are received, and transmit-only indicates transmitted data only. You can set the required attribute using **wlan ssid-profile refresh-direction command.** The receive-only mode does not use any null frames for refresh-direction.

The following procedure describes how to configure WLAN ageout refresh direction:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN > SSID**.
3. To edit an existing SSID profile, select the SSID profile you want to edit. To create a new SSID profile, click **+** and enter a name for the new SSID profile in the **Profile name** field.
4. Select either **RX-only** or **TX-only** from the **Station Refresh Direction** drop-down list.

   The Station Refresh Direction uses **bidirectional** data frames by default.
5. Click **Submit**.
6. Click **Pending Changes.**
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI command configures WLAN ageout refresh direction.

```
(host) [mynode] (config) #wlan ssid-profile <profile-name> refresh-direction
<bidirectional / rx-only / tx-only>
```

Configuring the SSID Profile

The following procedure describes how to configure the SSID profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN > SSID**.
3. To edit an existing SSID profile, select the SSID profile you want to edit. To create a new SSID profile, click **+** and enter a name for the new SSID profile in the **Profile name** field.
4. Configure your SSID settings.

   The configuration parameters are described in Table 100.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Table 100: *SSID Profile Parameters*

Parameter	Description
Advanced	
SSID Enable	Click this check box to enable or disable the SSID. The SSID is enabled by default.
ESSID	Name that uniquely identifies a wireless network. The network name, or ESSID can be up to 32 ASCII characters, if it contains unicode, depending on the language, the maximum characters vary. For example, ESSID could be up to 10 Chinese characters. If the ESSID includes spaces, you must enclose it in quotation marks.
WPA Passphrase	Enter the WPA passphrase. ▪ If the encryption type is wpa2-psk-aes, enter one of WPA passphrase, WPA Hexkey, or MPSK passphrase. The MPSK passphrase requires MAC authentication against a ClearPass Policy Manager server. ▪ If WPA passphrase, WPA Hexkey, and MPSK passphrase are entered, the MPSK passphrase takes precedence and a client has to use the MPSK passphrase as received from the ClearPass Policy Manager server. The MPSK passphrase requires MAC authentication against a ClearPass Policy Manager server. ▪ If WPA passphrase and WPA Hexkey are configured, that is, the encryption type is not mpsk-aes, only WPA Hexkey is considered.
Encryption	Select one of the following encryption types:
xSec	Encryption and tunneling of Layer-2 traffic between the controller and wired or wireless clients, or between controllers. To use xSec encryption, you must use a RADIUS authentication server. For clients, you must install the Funk Odyssey client software. Requires installation of the xSec license. For xSec between managed devices, you must install an xSec license in each managed device.
enhanced open	Enhanced open encryption with or without PMK caching.
wpa3-sae-aes	WPA3 with AES encryption using Simultaneous Authentication of Equals(SAE).

Table 100: *SSID Profile Parameters*

Parameter	Description
wpa3-aes-ccm-128	WPA3 with AES CCM encryption and dynamic keys using 802.1X.
wpa3-cnsa	WPA3 with AES GCM-256 encryption using CNSA (192 bit).
wpa3-aes-gcm-256	WPA3 with AES GCM-256 encryption.
static-wep	WEP with static keys.
dynamic-wep	WEP with dynamic keys.
wpa-tkip	WPA with TKIP encryption and dynamic keys using 802.1X.
wpa-aes	WPA with AES encryption and dynamic keys using 802.1X.
wpa-psk-tkip	WPA with TKIP encryption using a preshared key.
wpa-psk-aes	WPA with AES encryption using a preshared key.
wpa2-aes	WPA2 with AES encryption and dynamic keys using 802.1X.
wpa2-psk-aes	WPA2 with AES encryption using a preshared key.
wpa2-psk-tkip	WPA2 with TKIP encryption using a preshared key.
wpa2-tkip	WPA2 with TKIP encryption and dynamic keys using 802.1X.
mpsk-aes	MPSK with AES encryption.
Opmode transition	Enable backward compatibility for enhanced-open/wpa3-sae-aes opmodes.
Enable Management Frame Protection	When selected, the SSID supports MFP-capable and traditional clients. Management Frame Protection can only be enabled on SSIDs that support WPA2.
Require Management Frame Protection	When selected, the SSID supports Management Frame Protection-capable clients only. Management Frame Protection can only be enabled on SSIDs that support WPA2.
DTIM Interval	Specifies the interval, in milliseconds, between the sending of DTIM in the beacon. This is the maximum number of beacon cycles before unacknowledged network broadcasts are flushed. When using wireless clients that employ power management features to sleep, the client must revive at least once during the DTIM period to receive broadcasts
802.11a Basic Rates	Select the set of supported 802.11a rates, in Mbps, that are advertised in beacon frames and probe responses.
802.11a Transmit Rates	Select the set of 802.11a rates at which the AP is allowed to send data. The actual transmit rate depends on what the client is able to handle, based on information sent at the time of association and on the current error/loss rate of the client.

Table 100: *SSID Profile Parameters*

Parameter	Description
802.11g Basic Rates	Select the set of supported 802.11b/g rates that are advertised in beacon frames and probe responses.
802.11g Transmit Rates	Select the set of 802.11b/g rates at which the AP is allowed to send data. The actual transmit rate depends on what the client is able to handle, based on information sent at the time of association and on the current error/loss rate of the client.
Station Ageout Time	Time, in seconds, that a client is allowed to remain idle before being aged out.
Station Refresh Direction	The refresh direction the SSID profile. Select **RX-only** to use Receive-only data frames for refresh direction. Select **TX-only** to use Transmit-only data frames for refresh direction. Default: bidirectional.
Max Transmit Attempts	Maximum number of retries allowed for the AP to send a frame.
RTS Threshold	Wireless clients transmitting frames larger than this threshold must issue RTS and wait for the AP to respond with CTS. This helps prevent mid-air collisions for wireless clients that are not within wireless peer range and cannot detect when other wireless clients are transmitting. Default: 2333
Short Preamble	Click this check box to enable or disable a short preamble for 802.11b/g radios. Network performance may be higher when short preamble is enabled. In mixed radio environments, some 802.11b wireless client stations may experience difficulty associating with the AP using short preamble. To use only long preamble, disable short preamble. Legacy client devices that use only long preamble generally can be updated to support short preamble.
Max Associations	Maximum number of wireless clients per radio for the SSID (subject to an AP limit of 255 clients per radio). Default: 64
Wireless Multimedia (WMM)	Enables or disables WMM, also known as IEEE 802.11e Enhanced Distribution Coordination Function. WMM provides prioritization of specific traffic relative to other traffic in the network.
Wireless Multimedia U-APSD (WMM-UAPSD) Powersave	Enable WMM UAPSD powersave.
WMM TSPEC Min Inactivity Interval	Specify the minimum inactivity time-out threshold of WMM traffic. This setting is useful in environments where low inactivity interval time-outs are advertised, which may cause unwanted timeouts. The supported range is 0-3,600,000 milliseconds, and the default value is 0 milliseconds.
DSCP mapping for WMM voice AC (0-63)	DSCP used to map WMM voice traffic. The supported range is 0-63.
DSCP mapping for WMM video AC (0-63)	Select the DSCP used to map WMM video traffic. The supported range is 0-63.
DSCP mapping for WMM best-effort AC (0-63)	Select the DSCP value used to map WMM best-effort traffic. The supported range is 0-63.

Table 100: *SSID Profile Parameters*

Parameter	Description
DSCP mapping for WMM background AC (0-63)	Select the DSCP used to map WMM background traffic. The supported range is 0-63.
Hide SSID	Select this check box to enable or disable the hiding of the SSID name in beacon frames. Note that hiding the SSID does very little to increase security.
Deny_Broadcast Probes	When a client sends a broadcast probe request frame to search for all available SSIDs, this option controls whether or not the system responds for this SSID. When enabled, no response is sent and clients have to know the SSID in order to associate to the SSID. When disabled, a probe response frame is sent for this SSID.
Local Probe Request Threshold (dB)	Enter the SNR threshold below which incoming probe requests will get ignored. The supported range of values is 0-100 dB. A value of 0 disables this feature.
Disable Probe Retry	Click this check box to enable or disable battery MAC level retries for probe response frames. By default this parameter is enabled, which mean that MAC level retries for probe response frames is disabled. **NOTE:** This parameter is not supported for 200 Series access points.
Battery Boost	Converts multicast traffic to unicast before delivery to the client, thus allowing you to set a longer DTIM interval. The longer interval keeps associated wireless clients from activating their radios for multicast indication and delivery, leaving them in power-save mode longer and thus lengthening battery life. This parameter requires the PEFNG license.
WEP Key 1	First static WEP key associated with the key index. Can be 10 or 26 hex characters in length.
WEP Key 2	Second static WEP key associated with the key index. Can be 10 or 26 hex characters in length.
WEP Key 3	Third Static WEP key associated with the key index. Can be 10 or 26 hex characters in length.
WEP Key 4	Fourth Static WEP key associated with the key index. Can be 10 or 26 hex characters in length.
WEP Transmit Key Index	Key index that specifies which static WEP key is to be used. Can be 1, 2, 3, or 4.
WPA Hexkey	WPA PSK.
WPA Passphrase	WPA passphrase with which to generate a PSK.
Maximum Transmit Failures	The AP assumes the client has left and should be deauthorized when the AP detects this number of consecutive frames were not delivered because the maximum retry threshold has been exceeded.
BC/MC Rate Optimization	Click this check box to enable or disable scanning of all active stations currently associated to an AP to select the lowest transmission rate for broadcast and multicast frames. This option only applies to broadcast and multicast data frames; 802.11 management frames are transmitted at the lowest configured rate.

Table 100: *SSID Profile Parameters*

Parameter	Description
	NOTE: Do not enable this parameter unless instructed to do so by your Aruba technical support representative.
Rate Optimization for delivering EAPOL frames	Click this check box to use a more conservative rate for more reliable delivery of EAPOL frames. Default: Enabled
Strict Spectralink Voice Protocol (SVP)	Click this check box to enable Strict SVP
802.11g Beacon Rate	Click this drop-down list to select the beacon rate for 802.11g (use for DAS only). Using this parameter in normal operation may cause connectivity problems.
802.11a Beacon Rate	Click this drop-down list to select the beacon rate for 802.11a (use for DAS only). Using this parameter in normal operation may cause connectivity problems.
Video Multicast Rate Optimization	When configured, the managed device chooses the rate for video multicast frames. You can configure MCS rates as well. MCS is an important setting because it provides for potentially greater throughput. **NOTE:** The following information displays the MCS rate if the **Short guard interval in 20 MHz mode** setting in **High-throughput SSID profile** is either enabled or disabled: MCS Streams 20 MHz 20 MHz SGI --- ------- ------ ---------- 0 1 6.5 7.2 1 1 13.0 14.4 2 1 19.5 21.7 3 1 26.0 28.9 4 1 39.0 43.3 5 1 52.0 57.8 6 1 58.5 65.0 7 1 65.0 72.2 8 2 13.0 14.4 9 2 26.0 28.9 10 2 39.0 43.3 11 2 52.0 57.8 12 2 78.0 86.7 13 2 104.0 115.6 14 2 117.0 130.0 15 2 130.0 144.4 **NOTE:** The MCS rates for video multicast are supported in all 802.11n -capable APs. This is not supported in 320 Series AP.
Advertise QBSS Load IE	Click this check box to enable the AP to advertise the QBSS load element. The element includes the following parameters that provide information on the traffic situation: ▪ **Station count:** The total number of stations associated to the QBSS. ▪ **Channel utilization:** The percentage of time (normalized to 255) the

Table 100: *SSID Profile Parameters*

Parameter	Description
	channel is sensed to be busy. The access point uses either the physical or the virtual carrier sense mechanism to sense a busy channel. ▪ **Available admission capacity:** The remaining amount of medium time (measured as number of 32us/s) available for a station via explicit admission control. The QAP uses these parameters to decide whether to accept an admission control request. A wireless station uses these parameters to choose the appropriate access points. **NOTE:** Ensure that WMM is enabled for legacy APs to advertise the QBSS load element. For 802.11n APs, ensure that either WMM or high throughput is enabled.
Advertise Location Info	When this option is enabled, APs broadcast their location within an IE carried in Beacon frames and Probe Response frames. The AP's latitude, longitude and altitude can be configured on the **Configuration > Wireless> AP Installation** page of the managed device WebUI, or using the **provision-ap** command in the managed device command-line interface.
Advertise AP Name	If this parameter is enabled, APs will broadcast the AP name configured by the **ap-name** command. Default: Disabled
Enforce user vlan for open stations	Select this option to restrict data traffic from open stations to the VLAN assigned to the user. This option is disabled by default.
Enable OKC	OKC is a similar technique, not defined by 802.11i, available for authentication between multiple APs in a network where those APs are under common administrative control. An Aruba deployment with multiple APs under the control of a single managed device is one such example. Using OKC, a station roaming to any AP in the network will not have to complete a full authentication exchange, but will instead just perform the 4-way handshake to establish transient encryption keys.
General	
Enable Agile Multiband (MBO)	Select this option to enable or disable Agile Multiband on the SSID profile.
Advertize Cellular Data Capability attribute of MBO	Select this option for the AP to advertize Cellular Data Capability (CDC) for MBO.

The following CLI command configures the SSID profile:

```
(config) #wlan ssid-profile <profile>
```

WLAN Authentication

The WLAN Wizard allows you to define the type of authentication used by clients associating to a WLAN. The WLAN wizard is the recommended method for defining WLAN settings, but advanced users can also define authentication settings manually via the AAA profile.

The following procedure describes how to configure WLAN authentication:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > AAA Profiles** tab.
2. From the **AAA Profiles** list, select **AAA**.
3. To edit an existing AAA profile, select the AAA profile you want to edit. To create a new AAA profile, click **+** and enter a name for the new AAA profile in the **Profile name** field.
4. Configure the AAA profile parameters described in Table 101.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Table 101: *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. 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. This feature requires the PEFNG license.
User idle timeout	Specify the idle timeout value for the client in seconds. A value of 0, deletes the user immediately after disassociation from the wireless network. Valid range is 30-15300 in multiples of 30 seconds.
RADIUS Interim Accounting	When this option is enabled, the RADIUS accounting feature allows the managed device to send Interim-Update messages with current user statistics to the server at regular intervals. This option is disabled by default, allowing the managed device to send only start and stop messages to the RADIUS accounting server.
User derivation rules	Click the 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.
Reauthenticate wired user on VLAN change	When a wired user moves across VLANs, a trigger is created to reauthenticate this user. The default value is 'Enabled'.

Parameter	Description
Device Type Classification	When you select this option, the managed device will parse user-agent strings and attempt to identify the type of device connecting to the AP. When the device type classification is enabled, the Global client table shown in the **Monitoring>Network > All WLAN Clients** window shows each client's device type, if that client device can be identified.
Enforce DHCP	When you select this option, clients must obtain an IP using DHCP before they are allowed to associate to an AP. Enable this option when you create a user rule that assigns a specific role or VLAN based upon the client device's type. For details, see Working with User-Derived VLANs on page 505. If a client is removed from the user table by the "Logon user lifetime" AAA timer, then that client will not be able to send traffic until it renews it's DHCP. Enforce DHCP is available on the managed device for APs configured for tunnel or decrypt-tunnel forwarding mode only.
PAN firewalls Integration	Requires IP mapping at Palo Alto Networks firewalls.
Open SSID RADIUS Accounting	Initiates RADIUS accounting as soon as the user associates to an Open SSID without any authentication. Do not enable this parameter for wired users. If enabled, the managed device sends RADIUS accounting packets for unauthenticated wired users.

The following CLI commands configure a AAA profile with user derivation rule:

```
(host) [md] (config) #aaa profile <profile_name>
(host) [md] (AAA Profile <profile_name>) #user-derivation-rules <rule_name>
```

AirMatch is the next generation radio resource management service introduced in ArubaOS 8.0.0.0 for devices in a Mobility Master-Managed Device topology. AirMatch provides RF network resource allocation with unprecedented quality. It analyzes the past 24 hours of RF network statistics and proactively optimizes the network for the next day. Any RF plan change is applied in the early morning to minimize client disruption and maximize the user experience. AirMatch can react to detrimental RF events, such as radar and high noise levels, to allow the network to manage sudden changes in the RF environment.

Stand-alone controllers support only the ARM and ClientMatch features, which use automatic, infrastructure-based controls to maximize client performance and enhance the stability and predictability of the Wi-Fi network.



AirMatch and ARM cannot be used together. A Mobility Master deployment that includes managed devices does not support Adaptive Radio Management.

ArubaOS now also provides IPv6 support for devices in a Mobility Master-Managed Device topology in AirMatch. Hence, you can configure mixed deployment of devices with IPv4 and IPv6 addresses.

This section describes the following topics:

RF Management for Mobility Master Deployments with Managed Devices

The following sections provide a general overview of the RF management used by a multi-managed device deployment managed by Mobility Master.

- AirMatch RF Management Overview on page 572
- ClientMatch Overview on page 574

The sections below describe the procedures to configure AirMatch and ClientMatch:

- Configuring AirMatch on page 577
- Configuring ClientMatch on page 579

RF Management for Deployments with a Stand-alone Controller

The following sections provide a general overview of the RF management used by stand-alone controllers:

- RF Management for Stand-alone Controller Deployments on page 580
- ClientMatch Overview on page 574

The sections below describe the steps to configure advanced ARM settings and troubleshoot common ARM issues:

- ARM Coverage and Interference Metrics on page 586
- Configuring ARM Profiles on page 588
- Troubleshooting ARM on page 591

AirMatch RF Management Overview

The AirMatch channel and EIRP optimization features deprecate the channel planning and EIRP optimization features in the legacy ARM feature. AirMatch is supported on Mobility Master only, while legacy ARM channel optimization and EIRP features continue to be supported by stand-alone controllers running ArubaOS.

AirMatch channel planning evens out channel distributions in any size of network, and in any subset of the contiguous network (as much as allowed by the network configuration, regulatory domain, and AP hardware capability). AirMatch also minimizes channel coupling, where adjacent radios are assigned to the same channel. The computing power of Mobility Master impacts channel distribution calculations, so channel coupling may occasionally be allowed in complex networks to keep the computing time practical.

AirMatch EIRP planning automatically considers the local density of the network to manage the APs' coverage and MCS operation, and optimizes EIRP changes across neighboring AP radios in order to offer users the best roaming experience.

Table 102 describes some of the differences between the channel and EIRP optimization features supported by ArubaOS AirMatch and ArubaOS ARM.

Table 102: *AirMatch and ARM in ArubaOS*

Features	AirMatch	ARM
Initial Release	ArubaOS 8.0.0.0	ArubaOS 2.x
Supported Topology	Mobility Master / Managed device	Stand-alone controller
Run Period	24 hours	As little as 5 minutes
RF information used	Past 24 hours of RF data	Instantaneous snapshot of the RF environment
Deployment Time	5 AM (by default) , or any time necessary **NOTE:** Starting with ArubaOS 8.1.0.0, the deployment time for each managed device is based upon the time zone configured for that device. In ArubaOS 8.0.x, the deployment time for all managed devices was based upon the time zone of the Mobility Master server.	Any time necessary
Computing Time	Depends upon network size	Less than 1 second
Optimization Scope	The entire RF network	Each individual AP

AirMatch Channel Assignments

Each AP in a Mobility Master deployment measures its RF environment for a five minute period, every 30 minutes by default. The AP then sends AMON messages about the radio feasibility to the managed device based on the AP hardware capability, radio and regulatory domain, and RF neighbors. The managed device forwards these messages to the Mobility Master. The Mobility Master adds this information to a database, computes an optimal solution, and deploys the latest RF plan by sending

updated settings to the APs. By default, this configuration update is sent to each device at 5 AM (as per the system clock for each managed device), but time of this configuration update can be modified via the AirMatch profile.



An exception to this daily update is an automatic channel change due to a radar detection event or high noise interference. If an AP detects a radar event on its current operating channel, that AP automatically changes to another supported channel to avoid radar interference, and does not wait for the daily RF configuration update from the Mobility Master. An AP may also automatically change channels if a very high noise level is detected on the current channel, if at least one other channel is free of noise.

AirMatch moves a radio to a random channel when a radar event is detected, or if a high noise floor is detected on a non-static channel. AirMatch uses the criteria described in Table 103 to assign a new channel.

Table 103: *Channel Assignment Logic*

Issue Prompting Channel Change	Channel Selection Criteria
Detected radar	AirMatch selects a channel with a minimum interference index from the channels without high noise or a radar condition. If an AP detects radar on a channel, AirMatch will change the AP channel and will blacklist that channel for the next 24 hours.
High channel noise	If an AP detects high noise levels on a channel, AirMatch will change the AP channel and will blacklist that channel for the next 24 hours. The channel selection criteria varies between static and non-static channels. ■ If static channel is configured, the channel does not change due to a high noise condition. ■ For a non-static channel, AirMatch selects a channel with a minimum interference index from the channels without high noise or a radar condition.

Channel Quality Improvement Thresholds

The AirMatch channel quality improvement threshold allows you to select the minimum channel improvement that can trigger a new scheduled channel solution. The default threshold value is a 15% improvement. If a proposed channel change will not produce an improvement that meets or exceeds this threshold, AirMatch will not trigger a channel change.



This channel quality setting only applies to scheduled updates. If you manually trigger an update using the **airmatch runnow** command, AirMatch will deploy the new solution regardless of the level of improvement.

Initial RF Calculations

The database for the AirMatch service is empty when Mobility Master first boots up. When Mobility Master first detects APs on the network, it enters its initial optimization phase, collects data from all the APs, and generates an incremental solution every 30 minutes (by default) for the next eight hours. When this initial eight-hour period has elapsed, the AirMatch service will periodically calculate a new RF configuration for these devices.

When a new AP is deployed on a network with an active Mobility Master during the initial 8-hour AirMatch optimization phase that AP joins the network with its preassigned channel and transmission

power values. The AirMatch service detects the newly deployed AP on the network, restarts its RF computations, and sends an incremental RF configuration update to the new AP 30 minutes later. APs added to the network after the initial 8-hour optimization period will not receive an additional RF configuration update until the next scheduled update period.

Analytics Integration with Airmatch

Network Analytics Engine aggregates the network data, monitoring, and troubleshooting content of a particular network. Analytics and data mining systems use this content through APIs to detect and analyze problems. Starting from ArubaOS 8.4.0.0, the ArubaOS Mobility Master infrastructure is integrated with analytics engine. The analytics engine can push radio profile EIRP recommendations, channel-bandwidth recommendations, and regulatory domain profile recommendations to an AP.

For EIRP recommendations, AirMatch sees the recommended configurations from the analytics engine and computes EIRP min (eirp-min), EIRP max (eirp-max), and EIRP offset (eirp-offset). ArubaOS validates these recommendations and overrides the current configuration of an AP.



This feature is not supported in stand-alone mode or Mobility Controller mode.

ClientMatch Overview

ClientMatch continually monitors the RF neighborhood for each client to provide ongoing client band steering and load balancing, and enhanced AP reassignment for roaming mobile clients.



Legacy 802.11a/b/g devices do not support ClientMatch. When you enable ClientMatch on 802.11n-capable devices, ClientMatch overrides any settings configured for the legacy bandsteering or load balancing features. 802.11ac-capable devices do not support the legacy bandsteering, station hand off or load balancing settings, so these APs must be managed on using ClientMatch.

The managed device aggregates information it receives from all APs using ClientMatch, and maintains information for all associated clients in a database. The managed device shares this database with the APs (for their associated clients), and the APs use the information to compute the client-based RF neighborhood and determine which APs should be considered candidate APs for each client. When the managed device receives a client steer request from an AP, the managed device identifies the optimal AP candidate and manages the client's relocation to the desired radio. This is an improvement from previous releases, where ARM was managed exclusively by APs, without the larger perspective of the client RF neighborhood.

In Mobility Master / managed device deployments where APs are connected to a managed device that is associated to Mobility Master, the AP sends RF neighborhood information to the managed device, which then forwards that information to the Mobility Master. The Mobility Master receives probe reports from all managed devices and generates a VBR for each client. These VBRs are sent from the Mobility Master to the managed device, and then to the AP to which the client is associated. APs associated to a stand-alone controller receive and collect information about clients in their neighborhood, and periodically send this information to the controller, which in turn generates VBRs and sends them directly back to the APs.

ArubaOS now provides native IPv6 support in ClientMatch that allows band steering and sticky-client features in both pure IPv6 as well as dual-stack deployments. The native IPv6 support also ensures that the VBR for each client displays IPv6 address of the APs.

The following client or AP mismatch conditions are managed by ClientMatch:

- **Load Balancing**: ClientMatch balances clients across APs on different channels, based upon the client load on the APs and the SNR levels that the client detects from an underused AP. If an AP radio can support additional clients, the AP will participate in ClientMatch load balancing, and clients can be directed to that AP radio, subject to predefined SNR thresholds.
- **Sticky Clients**: ClientMatch also helps mobile clients that tend to stay associated to an AP despite low signal levels. APs using ClientMatch continually monitor the client RSSI as it roams between APs, and moves the client to an AP when a better radio match is found. This prevents mobile clients from remaining associated to APs with a less than ideal RSSI, which can cause poor connectivity and reduce performance for other clients associated with that AP.
- **Band Steering/Band Balancing**: APs using the ClientMatch feature monitor the RSSI for clients that advertise dual-band capability. If a client is currently associated to a 2.4 GHz radio, and the AP detects that the client has a good RSSI from the 5 GHz radio, the managed device attempts to steer the client to the 5 GHz radio, as long as the 5 GHz RSSI is not significantly worse than the 2.4 GHz RSSI, and the AP retains a suitable distribution of clients on each of its radios.
- **HE Steering**: 802.11ax clients are best compatible with 802.11ax capable radios, resulting in better throughput and spectral efficiency. When an 802.11ax client is associated with a lower radio, ClientMatch pushes the client to the best compatible 802.11ax radio for advanced capabilities. Though STA is in good health, and is 802.11ax capable, it still sometimes connects to lower radios. ClientMatch finds a potential 802.11ax radio on the same band and the client moves to the new 802.11ax radio.

This section describes the following topics:

- Incremental Rules-Based ClientMatch Updates on page 575
- BSS Transition Management Support on page 575
- Multi-Media Sync-Up on page 576
- Multi-User MIMO Steering on page 576
- Removing VBR Dependency on Probe Requests on page 576
- ClientInsight on page 577

Incremental Rules-Based ClientMatch Updates

The ClientMatch rules that manage client associations are based primarily upon the client RF environment and apply uniformly to all client types, regardless of device type or operating system. ArubaOS 8.0.0.0 supports incremental updates to ClientMatch rules to support network devices running newer operating systems that may be incompatible with the existing ClientMatch client association rules. This feature allows the managed device to use a newer set of ClientMatch rules without updating the entire operating system, reducing network downtime.

BSS Transition Management Support

The BSS Transition Management Support feature allows ClientMatch to steer devices using 802.11v BSS transition management standards for continuous wireless connectivity. This feature provides a seamless standards-compatible method of device steering in wireless networks, as 802.11v BSS transition management support has become increasingly common in wireless devices.

When ClientMatch attempts to steer the client to a more optimal AP, it sends out an 802.11v BSS transition management request to the 11v capable station and waits for a response.

1. ClientMatch begins a timeout session for the BSS transition management response or new association request to the desired AP.

2. If the request is rejected or the timeout session expires, ClientMatch is notified of the failed attempt and reinitiates the steer using the 802.11v BSS transition management request.

- If the client steer fails the maximum number of timeout attempts (default: 5), ClientMatch marks the client as 11v unsupported and falls back to using deauths to steer.
- If the client steer fails due to request rejection, ClientMatch does not mark the client as 11v unsupported and continues to attempt to steer using the 802.11v BSS transition management request.

Multi-Media Sync-Up

ClientMatch offers a tighter integration with multiple media-aware ALGs to provide better call quality for programs like Skype for Business (Skype4b) and Facetime. With ClientMatch's ability to understand various media protocols, clients are not steered to different APs in the middle of an active media session.

When a client participates in a call, the managed device learns about the media session and sends this information to the AP to which the client is currently associated, as part of the variable bitrate update. When the AP learns that the client is in a call, it will not attempt to steer the client to another AP until the managed device indicates that the call has ended, allowing calls to run more smoothly without any disruptions to the ongoing media flow.

Multi-User MIMO Steering

Multi-user MIMO, or MU-MIMO Steering, groups multi-user-capable (MU-capable) clients to maximize the likelihood of MIMO transmissions, which increases downstream throughput performance in 802.11ac Wave 2 (gen 2) APs. MU-MIMO runs on MU-capable clients with traffic flows and PHY channels compatible for multi-user transmissions. ClientMatch steers and aligns MU-MIMO-capable clients with MU-MIMO-capable radios using SNR values. Multiple MU-MIMO-capable clients can be grouped together on a MU-MIMO-capable radio.

Successful MU-MIMO transmissions depend on the following:

- Traffic streams that can be multiplexed for MIMO transmissions. This is dependent on packet length and traffic flow rates (packet arrival rates) from APs to the devices.
- MU-MIMO-capable clients associated to the same radio, whose PHY channel matrices are compatible for simultaneous multi-user transmissions

In an 802.11ac AP deployment, clients indicate VHT capabilities for probe requests and association requests, including MU-MIMO support. The APs and managed devices use this information to determine whether the client is MU-MIMO-capable.

After the MU-MIMO-capable clients are located, they are steered to an appropriate MU-MIMO-capable radio. MU-MIMO Steering ensures that steers are compatible with existing trigger thresholds, such as sticky clients and load-balancing. The multi-user SNR threshold of the target radio must be greater than the sticky client SNR threshold, and radios that exceed the client threshold are avoided to prevent the need for load-balancing.

Removing VBR Dependency on Probe Requests

ClientMatch has shifted its dependency on probe requests to the AM data feeds for virtual beacon report data. Instead of relying solely on client background scans during probe requests, which can cause limitations due to low scanning frequency, ClientMatch uses AM data feeds to gain more continuous, comprehensive client RSSI feeds. Along with probe requests, Air Monitor data feeds collect client information during AP scanning using the following frames:

- Block ACK
- Management frames
- NULL data frames
- Data frames with rates no higher than 36Mbps
- Control frames

ClientInsight

ClientInsight is a new feature that integrates ClientMatch with data analytics and insights from NetInsight. ClientInsight is designed to support the next generation data-driven wireless network automation.

Aruba NetInsight delivers network assurance by arming IT organizations with machine learning based analytics for proactively running today's fast paced networks. With automated insights and prescriptive recommendations, businesses can continuously adapt and improve the quality of experience for users and the Internet of Things (IoT).

ClientInsight allows for NetInsight to automatically create customized rules based on observations on ClientMatch outcomes steered during deployments. You can configure this feature using the NetInsight API. For further details about ClientInsight, refer to the Aruba NetInsight User Guide.

Datapath Health Monitoring with AMON

ArubaOS now enables enhanced visibility of an application or controller health using datapath and session information. The data and counters maintained by datapath can be helpful in assessing datapath CPU utilization and controller health.

NetInsight currently uses the **show datapath** command, and **AMON_HWMON_SYS_INFO_MESSAGE (71)** and **AMON_AP_SYSTEM_STATS(18) AMON** messages to indicate controller CPU and AP CPU utilization respectively. Starting from ArubaOS 8.5.0.0, the following list of new AMON messages are introduced to determine and assess the datapath CPU utilization and health.

- AMON_SOS_RES_UTIL_MESSAGE
- AMON_SOS_CPU_UTIL_STATS_MESSAGE
- AMON_SOS_DEBUG_DMA_MESSAGE
- AMON_SOS_BWM_MESSAGE
- AMON_SOS_MAINT_CNTR_MESSAGE
- AMON_SOS_CNTR_DESC_MESSAGE
- AMON_SOS_CNTR_VAL_MESSAGE



AMON_SOS_CNTR_DESC_MESSAGE and AMON_SOS_CNTR_VAL_MESSAGE messages are in Protobuf format and not the traditional AMON message format.

Configuring AirMatch

The range of RF settings that can be assigned to an the AP via the AirMatch feature is defined in the 2.4 GHz and 5 GHz radio profiles on the managed device. You can access these settings on the Mobility Master WebUI by selecting the configuration for the managed device from the configuration hierarchy, then navigating to the **Configuration > AP Groups** and **Configuration > Access Points** pages. Use these pages to specify the radio mode and range of channels and maximum channel bandwidth that can be assigned to an AP or AP group via an AirMatch solution. The AirMatch feature will not assign an

AP a channel that does not fall within the group of valid channels or channel bandwidth ranges allowed by that 2.4 GHz and 5 GHz radio profile used by that AP.

The AirMatch feature performs automatic daily updates by default, but you can use the Mobility Master WebUI or command-line interface to disable daily updates for APs at one or more configuration nodes, allowing those APs and retaining their existing RF configuration. If the AirMatch updates are changed from the default **enabled** setting to **disabled**, the Mobility Master continues to receive RF updates from the APs but Mobility Master does not execute any channel or EIRP changes.



- The AirMatch **disabled** setting is different from the ARM **disable** or **maintain** setting on a standalone controller. The ARM **disable** setting changes the AP radio channel and EIRP values back to the default values specified in the 2.4 GHz and 5 GHz radio profiles for that radio. The ARM **maintain** setting freezes the current radio channel and EIRP settings. In contrast, if you use AirMatch in a Mobility Master/Managed Device topology, the AirMatch **disabled** option simply means the centralized algorithm will stop selecting a new channel, bandwidth, or EIRP setting; the network operator still can override the previous settings assigned by AirMatch with static channel or EIRP values, and the AP radio can continue to voluntarily change channels to avoid radar interference or high noise levels.
- AirMatch supports manual dual 5 GHz mode selection in AP-344 access points and auto dual 5 GHz mode selection in AP-345 access points.

The following procedure describes how to define the most commonly used AirMatch configuration settings, but some advanced AirMatch settings are only available in the CLI. The following steps hold the existing AirMatch RF configuration and will disable future updates in ArubaOS 8.0.1.0 or later:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > Services > AirMatch**.
2. Click the **Automatically deploy AirMatch optimizations** toggle switch to enable this setting.
3. To change the time of the daily AirMatch RF updates, click the **Deploy daily at** drop-down list and select an update interval (in 24-hour format).
4. Click **Submit**.
5. Select **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.



In ArubaOS 8.0.0.0, the AirMatch WebUI was available at **Configuration > Services > More > AirMatch** page of the ArubaOS WebUI.

The following CLI command holds the existing AirMatch RF configuration and disable future updates in ArubaOS 8.0.1.0 or later:

```
(host) [mynode] (config) #airmatch profile schedule disable
```

The following CLI commands changes the time of the daily AirMatch RF updates from the default 5 AM to 2 AM, in the time zone of each managed device:

```
(host) [mynode] (config) #airmatch profile deploy-hour 2
```

Use the **quality-threshold** parameter to change the percentage of channel quality improvement that triggers a scheduled AirMatch RF update. If a proposed channel change will not produce an improvement that meets or exceeds this threshold, AirMatch does not trigger a channel plan.

```
(host) [mynode] (config) #airmatch profile quality-threshold <quality-threshold>
```



If scheduled updates are enabled, the new channel plan is deployed on the specified deployment hour only if it is improved by greater than this threshold value. A new EIRP plan is deployed on the deployment hour every day.

Use the Mobility Master command-line interface to manually initiate AirMatch RF computations and solution deployment instead of waiting for the next scheduled update period. Access the CLI and issue the following command:

```
(host) [mynode] #airmatch runnow full
```

The **airmatch ap freeze** command deploys the specified channel and EIRP values to a radio immediately, then freezes those values, regardless of whether the AirMatch RF planning feature is set to **enable** or **disable** mode. A radio set with the **airmatch ap freeze** command uses a static radio configuration until those settings get explicitly canceled with the **airmatch ap unfreeze** command. This command can be used to freeze either the channel or the EIRP value, or both values. For example, you can freeze the channel on an AP radio, while allowing the EIRP values to be updated by AirMatch.

```
(host)[mynode](config)# airmatch ap freeze {ip-addr <ip-addr>}|{ip6-addr <ip6-addr>}|{ap-
name <ap-name>}|{ap-group <ap-group>}|{all-aps} {band <band>}|{channel <channel>}|{eirp
<eirp>}{lms {lms-ip <lms-ip>}|{{lms-ipv6 <lms-ipv6>}}
```



Starting with ArubaOS 8.2.0.0, the **eirp** parameter supports the configuration of EIRP values in .1 dBm increments. 270 Series access points support both positive and negative EIRP values; all other APs support positive values only.

Unfreezing a radio configuration with the **airmatch ap unfreeze** command does not mean that there will automatically be an immediate change in the AP radio channel and EIRP values. It does, however, mean that the AirMatch algorithm can assign a new set of values at the next update.

```
(host)[mynode](config)# airmatch ap unfreeze {ip-addr <ip-addr>}|{ip6-addr <ip6-addr>}|
{ap-name <ap-name>}|{ap-group <ap-group>}|{all-aps} band <band> {channel <channel>}|{eirp
<eirp>}{lms {lms-ip <lms-ip>}|{{lms-ipv6 <lms-ipv6>}}
```

By default, each AP in a Mobility Master deployment measures its RF environment for a five minute duration, every 30 minutes by default. Mobility Master uses this information to compute an optimal solution, then deploys the latest RF plan by sending updated settings to the APs. Use the **ap system profile** command to modify these default report intervals, or to disable AirMatch reports to the APs.

```
(host) [mynode] (config) #ap system-profile <profile>
  airmatch-measure-duration <airmatch-measure-duration>
  airmatch-report-enabled
  airmatch-report-period <airmatch-report-period>
```

Use the **eirp-offset** parameter in the 2.4 Ghz and 5 Ghz radio profiles to manually adjust EIRP levels by defining an additional EIRP offset value (from -6 to 6 dB) that will be added to the AirMatch solution.

```
(host) [mynode] (config) #rf dot11a-radio-profile default eirp-offset 2
(host) [mynode] (config) #rf dot11g-radio-profile default eirp-offset 2
```

Use the **minimum-channel-bandwidth** parameter in the 2.4 Ghz and 5 Ghz radio profiles to set a specific channel bandwidth for APs associated to managed devices. This parameter varies from the maximum-channel-bandwidth parameter, in that the maximum-channel-bandwidth allows the AP to implement an AirMatch solution that may select all possible channel widths up to the selected maximum value.

```
(host) [mynode] (config) #rf dot11a-radio-profile default minimum-channel-bandwidth 80MHz
```

Configuring ClientMatch

Use the following procedures to disable or re-enable ClientMatch, and upload a Rules-Based ClientMatch update package.

Enabling and Disabling ClientMatch

ClientMatch is enabled by default. The procedure to disable and re-enable ClientMatch varies, depending upon whether your deployment consists of multiple managed devices managed by a Mobility Master, or whether your APs are all associated to a stand-alone controller.

Mobility Master Deployments

The following procedure describes how to enable or disable the 2.4 GHz and 5 GHz radio settings for an AP:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > AP Groups** page.
2. Select the AP group from the **AP Groups** table.
3. Click the **Radio** tab below the **AP Groups** table to display the AP radio settings.
4. Expand the **Client Control** section.
5. Click the **Client match** checkbox to enable or disable both 2.4 GHz and 5 GHz radio settings. Click **Submit.**
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

For more information on managing 2.4 GHz and 5 GHz radio settings, see 2.4 GHz and 5 GHz Radio RF Management on page 676.

Stand-alone Controller Deployments

For stand-alone controllers that do not have any associated managed devices, the ClientMatch feature is enabled and disabled in the ARM profile used by the AP, as described in Configuring ARM Profiles on page 588. Although default ClientMatch settings are recommended for most users, advanced ClientMatch settings can be configured using **rf arm-profile** commands in the command-line interface.

Uploading a Custom ClientMatch Rule Update Package

Use the WebUI or CLI to upload a custom update file of ClientMatch rules to the **/flash/config** folder on Mobility Master. This feature is not available for stand-alone controller deployments.

The following procedure describes how to upload a ClientMatch rule update package in ArubaOS 8.0.1.0 or later:

1. In the **Mobility Master** node hierarchy, select the device and navigate to **Diagnostics > Technical Support > Client Match Rules**.
2. Click **Upload File**, and then select a file to upload.
3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands upload a ClientMatch rule update package in ArubaOS 8.0.1.0 or later:

```
(host) [mynode] (config) #copy tftp: <tftphost> <filename> flash: <destname>
(host) [mynode] (config) #copy ftp: <ftphost> <user> <password> flash: <destname>
(host) [mynode] (config) #copy scp: <scphost> <username> <password> flash: <destname>
```

RF Management for Stand-alone Controller Deployments

Aruba's 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 Aruba AP in its current RF environment.

Aruba's ARM technology solves wireless networking challenges for stand-alone controllers in a large network deployment, dense deployment, or a network 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 continually monitors and adjusts radio resources to provide optimal network performance for APs associated to a stand-alone controller. 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.

This section describes the following features:

- ARM Monitoring and Management on page 581
- Traffic Shaping on page 584
- Cellular Handoff Assist

ARM Monitoring and Management

When ARM is enabled, the Aruba AP dynamically scans all 802.11 channels in its regulatory domain at regular intervals and will report everything it sees to the controller on each channel it scans (by default, 802.11n-capable APs scan channels in all regulatory domains). This includes, but is not limited to, data regarding WLAN coverage, interference, and intrusion detection. You can retrieve this information from the controller to get a quick health check of your WLAN deployment without having to walk around every part of a building with a network analyzer. For additional information on the individual matrix gathered on the AP's current assigned RF channel, see ARM Coverage and Interference Metrics on page 583.

This section describes the following topics:

- Maintaining Channel Quality on page 581
- Configuring ARM Scanning on page 581
- Understanding ARM Application Awareness on page 582
- Using Multi-Band ARM for 802.11a/802.11g Traffic on page 582
- 80 MHz Dynamic Bandwidth Switch on page 582
- ARM Coverage and Interference Metrics on page 583

Maintaining Channel Quality

Hybrid APs and Spectrum Monitors determine channel quality by measuring channel noise, non-Wi-Fi (interferer) utilization and duty-cycles, and certain types of retries. Regular APs using ARM derive channel quality values by measuring the noise floor for both 802.11 and non-802.11 noise on that channel.

The ARM algorithm is based on what the individual AP hears, so each AP on your WLAN can effectively "self heal" by compensating for changing scenarios like a broken antenna or blocked signals from neighboring APs. Additionally, ARM periodically collects information about neighboring APs to help each AP better adapt to its own changing environment.

Configuring ARM Scanning

The default ARM scanning interval is determined by the **scan-interval** parameter in the ARM profile. If the AP does not have any associated clients (or if most of its clients are inactive), ARM will dynamically readjust this default scan interval, allowing the AP to obtain better information about its RF neighborhood by scanning non-home channels more frequently. If an AP attempts to scan a non-home

channel but is unsuccessful, the AP will make additional attempts to rescan that channel before skipping it and continuing on to other channels.

The **Over the Air Updates** feature allows an AP to get information about its RF environment from its neighbors, even if the AP cannot scan. If you enable this feature, when an AP on the network scans a foreign (non-home) channel, it sends an Over-the-Air update in an 802.11 management frame that contains information about that home channel for that APl, the current transmission EIRP value of the home channel, and one-hop neighbors seen by that AP.

If ARM reports a high noise floor on a channel within a 40 MHz channel pair or 80 MHz channel set, ARM performs an additional 20 MHz scan on each channel within that channel pair or set, to determine the actual noise floor of each affected channel. This allows ARM to avoid assigning the overused channel, while still allowing channel assignments to the other unaffected channels in that channel pair or set.

Understanding ARM Application Awareness

Aruba 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 5 GHz band on dual-band APs. This frees up resources on the 2.4 GHz 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.

Using Multi-Band ARM for 802.11a/802.11g Traffic

It is recommended that you use the **multi-band** ARM assignment and **Mode Aware** ARM feature for single-radio APs in networks with traffic in the 802.11a and 802.11g bands. This feature allows a single-radio AP to dynamically change its radio bands based on current coverage on the configured band. This feature is enabled via the ARM profile used by the AP.

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. 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.

80 MHz Dynamic Bandwidth Switch

If an AP radio uses an 80 MHz channel, the radio only sends out frames when the entire 80 MHz channel is clear, even if the AP is sending only a 20 MHz management frame or 40 MHz data frame. As a result, throughput on the selected 80 MHz channel can be negatively impacted if interference occurs on both 20 MHz channels of the secondary 40 MHz channel.

The ARM dynamic bandwidth switch feature allows ARM to detect the 20 Mhz interferes in this situation and potentially move the AP radio to another 80 MHz channel, or change the AP transmissions to 40 MHz and use the primary 40 MHz channel instead.

When this feature is enabled, ARM starts a dynamic bandwidth switch observation window if load-aware scan rejects increase, *and* the clear channel assignment IBSS percentage (the percentage of channel traffic sent from that AP radio) drops below the value defined by the **dynamic-bw-cca-ibss-thresh** parameter.

If an observation window opens, and the clear channel assignment interference threshold exceeds the value defined by the **dynamic-bw-cca- intf-thresh** parameter, and the number of failed beacons from the radio exceeds the threshold defined by the **dynamic-bw-beacon- failed-thresh** parameter during that observation period, ARM will move the AP to another available 80 MHz channel with the minimum interference index. If no other 80 MHz channel is available, ARM downgrades the radio bandwidth to 40 MHz.



This feature is configured using the **rf arm-profile** command in the command-line interface. For more information refer to the *ArubaOS CLI Reference Guide*.

ARM Coverage and Interference Metrics

ARM computes coverage and interference metrics for each valid channel and chooses the best performing channel and transmit power settings for each AP's RF environment. Each AP gathers other metrics on their ARM-assigned channel to provide a snapshot of the current RF health state.

The following two metrics help the AP decide which channel and transmit power setting is best.

- **Coverage Index:** The AP uses this metric to measure RF coverage. The coverage index is calculated as x/y, where "x" is the AP's weighted calculation of the SNR on all valid APs on a specified 802.11 channel, and "y" is the weighted calculation of the AP 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

Cellular Handoff Assist

Some dual-network-capable devices, such as mobile phones, prefer to connect to Wi-Fi networks and may remain associated to a Wi-Fi network even when they experience poor performance at the edge of the Wi-Fi coverage area. When both the ClientMatch and the cellular handoff assist features are enabled, the cellular handoff assist feature can help a dual-mode, 3G or 4G-capable Wi-Fi device, such as an iPhone, iPad, or Android client at the end of a Wi-Fi network, switch from Wi-Fi to an alternate 3G or 4G radio that provides better network access. This feature is supported by iOS and Android devices only.

This feature is enabled via the Virtual AP profile for an AP or AP group. For more information on Virtual AP profiles and other WLAN configuration settings, see Basic WLAN Configuration on page 536

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. During every sampling period, airtime is allocated to each client, giving it the opportunity to 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.

Traffic shaping is configured in a traffic management profile.

The following procedure describes how to configure traffic shaping:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. From the **All Profiles** list, expand **QoS**.
3. Select the **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, click **+** and enter a name for a new profile.

The following table describes configuration settings available in the **General** and **Advanced** sections of the traffic management profile.

Table 104: *Traffic Management Profile Parameters*

Parameter	Description
General Settings	
Station Shaping Policy	Define Station Shaping Policy This feature has the following three options: ▪ **default-access:** Traffic shaping is disabled, and client performance is dependent on MAC contention resolution. This is the default traffic shaping setting. ▪ **fair-access:** Each client gets the same airtime, regardless of client capability and capacity. This option is useful in environments like a training facility or exam hall, where a mix of 802.11a/g, 802.11g, and 802.11n clients need equal to network resources, regardless of their capabilities. The bw-alloc parameter of a traffic management profile allows you to set a minimum bandwidth to be allocated to a virtual AP profile when there is congestion on the wireless network. You must set traffic shaping to fair-access to use this bandwidth allocation value for an individual virtual AP. ▪ **preferred-access:** High-throughput (802.11n) clients do not get penalized because of slower 802.11a/g or 802.11b transmissions that take more air time due to lower rates. Similarly, faster 802.11a/g clients get more access than 802.11b clients.
Advanced Settings	
Proportional BW Allocation	You can allocate a maximum bandwidth, as a percentage of available bandwidth to a virtual AP. To assign a percentage of bandwidth to a virtual AP: 1. Click **+** in the **Proportional BW Allocation** table. 2. Select the VAP profile to which you would like to allocate a bandwidth share from the **virtual_ap** drop-down list. 3. Specify the percentage of bandwidth to be allocated to the VAP in the **share** field. 4. Click the **hard_limit** drop-down list and select the mode for restricting the bandwidth for the VAP. Select the **soft** limit check box if you want to restrict the bandwidth for this VAP when there is a congestion on the wireless network. If you do want to restrict the bandwidth even when here there is congestion, select the **hard** limit option. 5. Click **OK**. 6. Repeat steps 1-5 to assign any remaining bandwidth to additional VAPs, if desired. To remove a VAP from the list of VAPs with allocated bandwidth, select the VAP from the **Proportional BW Allocation** table and click **Delete**.
Report Interval	Number of minutes between bandwidth usage reports. Range: 1-99 minutes Default value is 5 minutes.

5. Click **Save**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI command configures and enables traffic shaping:

```
(host)[mynode](config) #wlan traffic-management-profile <profile> shaping-policy default-access|fair-access|preferred-access
```

The following CLI command applies an 802.11a or 802.11g traffic management profile to an AP group or an individual AP.

```
(host) [mynode] (config) #ap-group <name> dot11a-traffic-mgmt-profile|dot11g-traffic-mgmt-profile <profile>
```

802.11ad

IEEE 802.11ad (WiGig) is a multi-gigabit Wi-Fi technology that allows managed devices to communicate at multi-gigabit speeds over a 60 GHz band. This technology comprises two radios, 5 GHz and 60 GHz.

The 802.11ad technology is supported by mesh networks and its radios support the following 2 GHz channels:

1. 1 (58320MHz)
2. 2 (60480MHz)
3. 3 (62640MHz)
4. 4 (64800MHz)

Managed devices can add the 802.11ad radio profile to store related configurations and to add 60 GHz channel entries to the regulatory domain profile.

ARM Coverage and Interference Metrics

ARM computes coverage and interference metrics for each valid channel, and chooses the best performing channel and transmit power settings for each AP's RF environment. Each AP gathers other metrics on their ARM-assigned channel to provide a snapshot of the current RF health state.

The following two metrics help the AP decide which channel and transmit power setting is best:

- **Coverage Index:** The AP uses this metric to measure RF coverage. The coverage index is calculated as x/y, where "x" is the AP's weighted calculation of the SNR on all valid APs on a specified 802.11 channel, and "y" is the weighted calculation of the AP SNR the neighboring APs see on that channel.
- To view these values for an AP in your current WLAN environment, issue the **show ap arm rf-summary ap-name <ap-name>** CLI command, 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. This metric is calculated and weighted for all APs, including third-party APs seen on a specified 802.11 channel and its adjacent channels. An adjacent channel is 25 MHz of the primary 20 MHz channel. For example, channel 40 is adjacent, because it is 20 MHz away from channel 36. 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 **show ap arm rf-summary ap-name <ap-name>** CLI command, then add the values *a+b+c+d.*

The following CLI example describes the interference index calculation involving four metric values:

```
(host) [md] (config) # show ap arm rf-summary ap-name ap1
Channel Summary
---------------
channel  retry  phy-err  mac-err  noise  util(Qual)    cov-idx(Total)  intf_idx
(Total)
-------  -----  -------  -------  -----  ----------    --------------  ------------
---
149      0      0        0        96     8/7/0/0/99    8/0(8)          18/11//4/1
(34)
153      0      0        0        92     6/3/2/0/99    3/0(3)          32/22//3/0
(57)
157      0      0        0        92     3/0/2/0/99    7/0(7)          50/13//2/5
(70)
161      0      0        0        92     3/0/2/0/95    0/1(1)          7/17//5/0
(29)
```

For channel 149, the interference index values are 18/11//4/1 where:

a= 18 (due to neighbors on channel 149)

b = 11 (due to neighbors on channel 153)

c = 4 (due to two-hop neighbors on channel 149)

d = 1 (due to two-hop neighbors on channel 153)

Hence, the interference index total on channel 149 is a+b+c+d i.e., 18+11+4+1 = 34.

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 **show ap arm rf-summary ip-addr <ap ip address>** or **show ap arm rf-summary ap-name <ap-name>** CLI command.

- Amount of Retry 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

The following enhancements are introduced in ArubaOS 8.0.0.0 to resolve issues that occur with the distributed channel/power algorithm:

- **Push random channel assignments to APs:** To support the random channel assignment feature, set the **Assignment** parameter in the ARM profile to **maintain**. Once this is done, random channels are pushed from the managed device STM/SAPM to APs that belong to a specific ap-group. This helps in replacing the dynamic channel change solution in a high density environment, thereby overcoming the issue with convergence. Random channel assignment helps in certain customer deployments where administrators want to control channels assigned and also for initial channel assignment to seed ARM channel computation.
- **Reduce interference channel change:** To reduce the number of interference channel changes and to configure the weight of interfering APs when calculating the interference index, the **interfering-ap-weight** parameter is introduced in the **rf-arm-profile** command. Before this enhancement was introduced, the value of the interfering AP (uncontrollable AP) was similar to the valid AP (controllable AP).

Configuring ARM Profiles

ARM profile settings are divided into two categories: **General, Scanning** and **Advanced.** The general ARM settings include general configuration parameters such as channel and power assignments and minimum and maximum allowed EIRP values.



Most network environments do not require any changes to the **Scanning** or **Advanced** categories of ARM configuration settings. If, however, your network supports a large amount of VoIP or Video traffic, or if you have unusually high security requirements you may want to manually adjust the basic ARM thresholds.

Default Profiles

When you create a new AP group and modify any of the ARM settings for that group, ArubaOS creates a unique profile for that AP group. The settings in these default profiles may vary, depending upon the radio type. The default ARM profile for a 2.4 GHz radio is Default-g, and the default profile for a 5 GHz radio is Default-a.

This section describes how to manually configure an ARM profile.

Manually Configuring an ARM Profile

The range of RF settings that can be assigned to an AP are defined in the 2.4 GHz and 5 GHz radio profiles. You can access these settings on the Mobility Master WebUI by selecting the configuration for the managed device from the configuration hierarchy, then navigating to the **Configuration > AP Groups > Radio** page. However, advanced ARM settings can be edited using the WebUI or CLI.



The ARM profile also includes advanced ClientMatch settings that can be configured through the command-line interface only. The default values for these settings are recommended for most users, and caution should be used when changing them to a non-default value. For complete details on all ClientMatch configuration settings, refer to the *ArubaOS CLI Reference Guide*.

The following procedure describes how to configure an ARM profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. Select **RF Management** from the **All Profiles** list, and then click **Adaptive Radio Management (ARM)**.
3. Select the ARM profile you want to edit, or create a new profile by clicking **+** and entering a name for the new profile in the **Profile Name** field. The ARM profile settings are divided into three sections, **General, Scanning** and **Advanced.** The profile parameters in each section are described in Table 105.

Table 105: *ARM Profile Configuration Parameters*

Parameter	Description	Default
General		
ClientMatch	The ClientMatch feature helps optimize network resources by balancing clients across channels, regardless of whether the AP or the managed device is responding to the wireless clients' probe requests.	Enabled

Table 105: *ARM Profile Configuration Parameters*

Parameter	Description	Default
	If enabled, the managed device compares whether an AP has more clients than its neighboring APs on other channels. If an AP's client load is at or over a predetermined threshold as compared to its immediate neighbors, or if a neighboring Aruba AP on another channel does not have any clients, load balancing will be enabled on that AP. This feature is enabled by default. For details, see ClientMatch Overview on page 574.	
Scanning		
Scanning	The **Scanning** check box enables or disables AP scanning across multiple channels. This check box is selected by default. Do not disable scanning unless you want to disable ARM and manually configure AP channel and transmission power. Disabling this option also disables the following scanning features: ■ Multi Band Scan ■ Rogue AP Aware ■ VoIP Aware Scan ■ Power Save Scan	Enabled
Multi Band Scan	If enabled, single radio channel APs scan for rogue APs across multiple channels. This option requires that **Scanning** is also enabled. (The **Multi Band Scan** option does not apply to APs that have two radios, as these devices already scan across multiple channels. If one of these dual-radio devices are assigned an ARM profile with Multi Band enabled, that device will ignore this setting.)	Enabled
VoIP Aware Scan	Aruba's VoIP 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.	Enabled
Power Save Aware Scan	If enabled, the AP will not scan a different channel if it has one or more clients that is in power save mode.	Disabled
Video Aware Scan	As long as there is at least one video frame every 100 mSec the AP will reject an ARM scanning request. Note that for each radio interface, video frames must be defined in one of two ways: ■ Classify the frame as video traffic via a session ACL. ■ Enable WMM on the WLAN's SSID profile and define a specific DSCP value as a video stream. Next, create a session ACL to tag the video traffic with the DSCP value.	Enabled
Scan Mode	By default, 802.11n-capable APs scan channels within all regulatory domains. To limit the AP scans to just the regulatory domain for that AP, click the **Scan Mode** drop-down list and select **reg-domain**.	all-reg-domain

Table 105: *ARM Profile Configuration Parameters*

Parameter	Description	Default
	NOTE: This setting does not apply to APs that do not support 802.11n; these APs will scan their regulatory domain only.	
Advanced		
Rogue AP Aware	If you have enabled both the **Scanning** and **Rogue AP options**, Aruba APs may change channels to contain off-channel rogue APs with active clients. This security features allows APs to change channels even if the **Client Aware** setting is disabled. This setting is disabled by default, and should only be enabled in high-security environments where security requirements are allowed to consume higher levels of network resources. You may prefer to receive Rogue AP alerts via SNMP traps or syslog events.	Disabled
Active Scan	When you enable **Active Scan**, an AP initiates active scanning via probe request. This option elicits more information from nearby APs, but also creates additional management traffic on the network. **Active Scan** is disabled by default, and should *not be enabled* except under the direct supervision of Aruba Support.	Disabled
ARM Over the Air Updates	The **ARM Over the Air Updates** option allows an AP to get information about its RF environment from its neighbors, even the AP cannot scan. If this feature is enabled, when an AP on the network scans a foreign (non-home) channel, it sends other APs an Over-the-Air update in an 802.11 management frame that contains information about the home channel for the scanning AP, the current transmission EIRP value of its home channel, and one-hop neighbors seen by that AP. Default: enabled	Enabled
Channel Quality Threshold	Channel quality percentage below which ARM initiates a channel change. Range: 0-100%	70%
Channel Quality Wait Time	If channel quality is below the specified channel quality threshold for this wait time period, ARM initiates a channel change. Range:1-3600 seconds	120 seconds
Minimum Scan Time	Minimum number of times a channel must be scanned before it is considered for assignment. Range: 0–2,147,483,647 scans. It is recommended to use a **Minimum Scan Time** between 1–20 scans.	8 scans
Load Aware Scan Threshold	Load aware ARM preserves network resources during periods of high traffic by temporarily halting ARM scanning if the load for the AP gets too high. The **Load Aware Scan Threshold** is the traffic throughput level an AP must reach before it stops scanning. Range: 0–20,000,000 bytes/second. (Specify 0 to disable this feature.)	1250000 Bps

Table 105: *ARM Profile Configuration Parameters*

Parameter	Description	Default
High Noise Backoff Time	Duration in mins for blacklisting AirMatch Solver's channel after two consecutive high noise detections. Setting to 0 will disable the backoff window. Range: 0-3600 minutes	720 minutes
Radar Backoff Time	Duration in mins for blacklisting AirMatch Solver's channel after two consecutive radar detections. Setting to 0 will disable the backoff window. Range: 0-3600 minutes	720 minutes

The following CLI command creates a new ARM profile or modifies an existing profile:

```
(host) [mynode] (config) #rf arm-profile <profile>
```

Dynamic Bandwidth Switch

ARM's dynamic bandwidth switch feature provides capability for ARM to detect the 20 Mhz interferer by reading the CCA statistics and other radio statistics. Once the signatures are detected, ARM moves to another 80 Mhz channel or downgrades to 40 Mhz. This feature only works when **dynamic-bw** parameter is enabled and ARM is set to use 80 Mhz assignment.



If ARM decides to downgrade the bandwidth to 40 MHz, then it will upgrade back to 80 Mhz after the clear time based on the volume of the traffic.

Enabling Dynamic Bandwidth Switch

The following CLI commands enable and set dynamic bandwidth switch:

```
(host) [mynode] (config) #rf arm-profile default
(host) [mynode] (Adaptive Radio Management (ARM) profile "default") #dynamic-bw
(host) [mynode] (Adaptive Radio Management (ARM) profile "default") #dynamic-bw-beacon-
failed-thresh
(host) [mynode] (Adaptive Radio Management (ARM) profile "default") #dynamic-bw-cca-ibss-
thresh
(host) [mynode] (Adaptive Radio Management (ARM) profile "default") #dynamic-bw-cca-intf-
thresh
(host) [mynode] (Adaptive Radio Management (ARM) profile "default") #dynamic-bw-clear-
time
(host) [mynode] (Adaptive Radio Management (ARM) profile "default") #dynamic-bw-wait-time
```

Troubleshooting ARM

If ARM is enabled but does not seem to be working properly, try some of the troubleshooting tips below.

Too many APs on the Same Channel

If many APs are selecting the same RF channel, there may be excessive interference on the other valid 802.11 channels. Issue the CLI commands **show ap arm rf-summary ap-name <ap-name>** or **show ap arm rf-summary ip-addr <ap ip address>**and calculate the Interference index (*intf_idx*) for all the valid channels.

An AP will only move to a new channel if the new channel has a lower interference index value than the current channel. The ARM Free Channel Index parameter specifies the required difference between two

interference index values. If this value is set too high, the AP will not switch channels, even if the interference is slightly lower on another channel. Lower the Free Channel Index to improve the likelihood that the AP will switch to a better channel.

Wireless Clients Report a Low Signal Level

If APs detect strong signals from other APs on the same channel, they may decrease their power levels accordingly. Issue the CLI commands **show ap arm rf-summary ap-name <ap-name>** or **show ap arm rf-summary ip-addr <ap ip address>** for all APs and check their current coverage index (*cov-idx*). If the AP's coverage index is at or higher than the configured coverage index value, then the APs have correctly chosen the transmit power setting. To manually increase the minimum power level for the APs using a specific ARM profile, define a higher minimum value with the 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.

APs Detect Errors but Do Not Change Channels

First, ensure that ARM error checking is enabled. The ARM Error Rate Threshold should be set to a percentage higher than zero. The suggested configuration value for the ARM Error Rate Threshold is 30–50%.

APs Don't Change Channels Due to Channel Noise

APs will only change channels due to interference if you enable ARM noise checking. Check to verify that the ARM Noise Threshold is set to a value higher than 0 dBm. The suggested setting for this threshold is 75 dBm.

Chapter 26
Wireless Intrusion Prevention

The ArubaOS WIP features and configurations are discussed in this chapter. WIP offers a wide selection of intrusion detection and protection features to protect the network against wireless threats. The WIP configuration is done on Mobility Master in the network.

These features do not require an RFprotect license:

- Rogue AP classification techniques other than AP classification rules
- Rogue containment
- Wired containment
- Wireless containment without Tarpit

For details on commands refer to the *ArubaOS CLI Reference Guide*.

This chapter contains the following topics:

- Monitoring the Security Dashboard on page 593
- Detecting Rogue APs on page 594
- Working with Intrusion Detection on page 597
- Configuring Intrusion Protection on page 608
- Configuring the WLAN Management System on page 611
- Understanding Client Blacklisting on page 618
- Working with WIP Advanced Features on page 621
- Configuring TotalWatch on page 621
- Administering TotalWatch on page 624
- Tarpit Shielding Overview on page 625
- Configuring Tarpit Shielding on page 625

Monitoring the Security Dashboard

The **Security** page of **Dashboard** in the WebUI, allows you to monitor the detection and protection of wireless intrusions in your network.

The dashboard's sections — **Detected Radios, Detected Clients** and **Events**—contain data as links. When these links are clicked, they arrange, filter, and display the appropriate information in a new page.

For example, in the **Detected Radio** section, if you click a number on the number against **Rogue**, a table filters and arranges information about those classified Rogue APs in a new page. Use the scroll bar at the right to view all the Rogue APs.



The term events in this section is meant to include security threats, vulnerabilities, attacks (intrusion or Denial of Service) and other similarly related events.

Similarly, the **Event** section contains data links. You can click on these data links to view information, in a new page, related to the event you selected. You can use the scroll bar at the right to view all the events.

Detecting Rogue APs

The most important WIP functionality is the ability to classify an AP as a potential security threat. An AP is considered to be rogue if it is both unauthorized and plugged in to the wired side of the network. An AP is considered to be interfering 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.

This section describes the following topics:

- WIDS Containment Enhancements on page 594
- Understanding Classification Terminology on page 594
- Understanding Classification Methodology on page 595
- Understanding AP Classification Rules on page 596
- Understanding Rule Matching on page 597

WIDS Containment Enhancements

Air Monitor enabled APs detect and mitigate possible security threats in a wireless network. Air Monitor supports containment of rogue APs and prevents clients from associating with rogue APs. Air Monitor sends tarpit or deauthentication containment frames if any of the following criteria is met:

- When an AP is marked for DOS, a single broadcast deauthentication frame is sent for disassociation and if stations do not honor the broadcast message, two unicast deauthentication frames are sent to disassociate the station from the AP and vice versa.
- To disassociate a valid station from the non-valid AP, a unicast deauthentication frame is sent from the station's MAC address to the AP and vice versa.
- AP impersonation is active and it disassociates all stations from the invalid AP by sending unicast deauthentication frames.

Understanding Classification Terminology

APs and clients are discovered during scanning of the wireless medium, and they are classified into various groups. The AP and client classification definitions are in Table 106 and Table 107.

Table 106: *AP Classification Definition*

Classification	Description
Authorized	An AP that is part of the enterprise providing WLAN service.
Neighbor	A neighboring AP is when the BSSIDs are known. Once classified, a neighboring AP does not change its state.
Interfering	An AP that is seen in the RF environment but is not connected to the wired network. An interfering AP is not considered a direct security threat since it is not connected to the wired network. For example, an interfering AP can be an AP that belongs to a neighboring office's WLAN but is not part of your WLAN network.
Rogue	An unauthorized AP that is plugged into the wired side of the network.

Classification	Description
Suspected Rogue	A suspected rogue AP is an unauthorized AP that may be plugged into the wired side of the network.
Contained	An AP for which DoS is enabled manually.

Table 107: *Client Classification Definitions*

Classification	Description
Authorized	Any client that successfully authenticates with a valid AP and passes encrypted traffic.
Contained	Any clients for which DoS is enabled manually.
Interfering	A client associated to any AP and is not valid.

Understanding Classification Methodology

A discovered AP is classified as a rogue or a suspected rogue by the following methods:

- Internal heuristics
- AP classification rules
- Manually by the user

The internal heuristics works by checking if the discovered AP is communicating with a wired device on the customer network. This is done by matching the MAC address of devices that are on the discovered AP's network with that of the user's wired network. The MAC of the device on the discovered AP's network is known as the Match MAC. The ways in which the matching of wired MACs occurs is detailed in the sections Understanding Match Methods on page 595 and Understanding Match Types on page 596.

This section describes the following topics:

- Understanding Match Methods on page 595
- Understanding Match Types on page 596
- Understanding Suspected Rogue Confidence Level on page 596

Understanding Match Methods

The match methods are:

- Plus One—The match MAC matches a device whose MAC address' last bit was one more than that of the Match MAC.
- Minus One—The match MAC matches a device whose MAC address' last bit was one less than that of the Match MAC.
- Equal—The match was against the same MAC address.
- OUI—The match was against the manufacturer's OUI of the wired device.

The classification details for **Discovered Radios** and **Discovered clients** are available by clicking on their respective section icons in the **Dashboard > Security** page of the WebUI. The information is also available in the **show wms rogue-ap** command.

Understanding Match Types

- **Eth-Wired-MAC:** The MAC addresses of wired devices learned by an AP on its Ethernet interface.
- **GW-Wired-MAC:** The collection of Gateway MACs of all APs across Mobility Master and managed devices.
- **AP-Wired-MAC:** The MAC addresses of wired devices learned by monitoring traffic out of other valid and rogue APs.
- **Config-Wired-MAC:** The MAC addresses that are configured by the user, typically that of well-known servers in the network.
- **Manual:** User-triggered classification.
- **External-Wired-MAC:** The MAC address matched a set of known wired devices that are maintained in an external database.
- **Mobility-Manager:** The classification was determined by the mobility manager, AMP.
- **Classification-off:** AP is classified as rogue because classification has been disabled, causing all non-authorized APs to be classified as rogue.
- **Propagated-Wired-MAC:** The MAC addresses of wired devices learned by a different AP than the one that uses it for classifying a rogue.
- **Base-BSSID-Override:** The classification was derived from another BSSID, which belongs to the same AP that supports multiple BSSIDs on the radio interface. For Aruba OUIs, if the base BSSID of a beacon matches the base BSSID of a known valid BSSID then the new BSSID is not considered to be valid.
- **AP-Rule:** A user-defined AP classification rule has matched.
- **System-Wired-MAC:** The MAC addresses of wired devices learned on the managed device.
- **System-Gateway-MAC:** The Gateway MAC addresses learned on the managed device.

Understanding Suspected Rogue Confidence Level

A suspected rogue AP is a potential threat to the WLAN infrastructure. A suspected rogue AP has a confidence level associated with it. An AP can be marked as a suspected rogue if it is determined to be a potential threat on the wired network, or if it matches a user-defined classification rule.

The suspected-rogue classification mechanisms are:

- Each mechanism that causes a suspected-rogue classification is assigned a confidence level increment of 20%.
- AP classification rules have a configured confidence level.
- When a mechanism matches a previously unmatched mechanism, the confidence level increment associated with that mechanism is added to the current confidence level (the confidence level starts at zero).
- The confidence level is capped at 100%.
- If your managed device reboots, your suspected-rogue APs are not checked against any new rules that were configured after the reboot. Without this restriction, all the mechanisms that classified your APs as suspected-rogues may trigger again, causing the confidence level to surpass its cap of 100%. You can explicitly mark an AP as “interfering” to trigger all new rules to match against it.

Understanding AP Classification Rules

AP classification rule configuration is performed only on Mobility Master. If AMP is enabled via the **mobility-manager** command, then processing of the AP classification rules is disabled on Mobility Master. A rule is identified by its ASCII character string name (32 characters maximum). The AP classification rules have one of the following specifications:

- SSID of the AP
- SNR of the AP
- Discovered-AP-Count or the number of APs that can see the AP

This following topics provide information on understanding SSID, SNR, Discovered-AP-Count specifications, and sample rules:

Understanding SSID specification

Each rule can have up to 6 SSID parameters. If one or more SSIDs are specified in a rule, an option of whether to match any of the SSIDs or not match all of the SSIDs can be specified. The default is to check for a match operation.

Understanding SNR specification

Each rule can have only one specification of the SNR. A minimum and/or maximum can be specified in each rule, and the specification is in SNR (db).

Understanding Discovered-AP-Count specification

Each rule can have only one specification of the Discovered-AP-Count. Each rule can specify a minimum or maximum of the Discovered-AP-count. The minimum or maximum operation must be specified if the Discovered-AP-count is specified. The default setting is to check for the minimum discovered-AP-count.

Sample Rules

If SSID equals xyz AND SNR > 40 then classify AP as suspected-rogue with conf-level-increment of 20

If SNR > 60 and DISCOVERING_APS > 2, then classify AP as suspected-rogue with conf-level increment of 35

If SSID equals 'XYZ', then classify AP as known-neighbor

Understanding Rule Matching

A rule must be enabled before it is matched. A maximum of 32 rules can be created with a maximum of 16 rules simultaneously active. If a rule matches, an AP is classified either as **Suspected Rogue** or as **Neighbor**

For an AP classified as **Suspected Rogue**, an associated confidence-level is provided (minimum is 5%).

The following mechanism is used for rule matching:

- When all the conditions specified in the rule evaluate to true, the rule matches.
- If multiple rules match, causing the AP to be classified as a Suspected Rogue, the confidence level of each rule is aggregated to determine the confidence level of the classification.
- When multiple rules match and any one of those matching rules cause the AP to be classified as a Neighbor, then the AP is classified as Neighbor.
- APs classified as either Neighbor or Suspected Rogue will attempt to match any configured AP rule.
- Once a rule matches an AP, the same rule will not be checked for the AP.
- When the managed device reboots, no attempt to match a previously matched AP is made.
- If a rule is disabled or modified, all APs that were previously classified based on that rule will continue to be in the newly classified state.

Working with Intrusion Detection

This section covers Infrastructure and Client Intrusion Detections as described in the following topics:

- Understanding Infrastructure Intrusion Detection on page 598
- Understanding Client Intrusion Detection on page 604

Understanding Infrastructure Intrusion Detection

Detecting attacks against the infrastructure is critical in avoiding attacks that may lead to a large-scale DoS attack or a security breach. This group of features detects attacks against the WLAN infrastructure, which consists of authorized APs, the RF medium, and the wired network. An authorized or valid-AP is defined as an AP that belongs to the WLAN infrastructure. The AP is either an Aruba AP or a third party AP. ArubaOS automatically learns authorized Aruba APs.

Table 108 presents a summary of the intrusion infrastructure detection features with their related commands, traps, and syslog identification. Feature details follow the table.

Table 108: *Infrastructure Detection Summary*

Feature	Command	Trap	Syslog ID

Detecting an 802.11n 40 MHz Intolerance Setting on page 601	`ids dos-profile <profile-name>`
	`detect-ht-40mhz-intolerance`
	`client-ht-40mhz-intol-quiet-time`
		wlsxHT40MHzIntoleranceAP
		wlsxHT40MHzIntoleranceSta
			126052, 126053, 127052, 127053

Detecting Active 802.11n Greenfield Mode on page 601	`ids unauthorized-device-profile <profile-name>` `detect-ht-greenfield`	wlsxHtGreenfieldSupported	126054, 127054
Detecting Ad hoc Networks on page 601	`ids unauthorized-device-profile <profile-name>` `detect-adhoc-network`	wlsxNAdhocNetwork	126033, 127033

Detecting an Ad hoc Network Using a Valid SSID on page 601	`ids unauthorized-device-profile <profile-name>`
	`detect-adhoc-using-valid-ssid`
	`adhoc-using-valid-ssid-quiet-time`
		wlsxAdhocUsingValidSSID	126068, 127068

Detecting an AP Flood Attack on page 602	`ids dos-profile <profile-name>`
	`detect-ap-flood`
	`ap-flood-threshold`
	`ap-flood-inc-time`
	`ap-flood-quiet-time`
		wlsxApFloodAttack	126034, 127034

Detecting AP Impersonation on page 602	`ids impersonation-profile <profile-name>`
	`detect-ap-impersonation`
	`beacon-diff-threshold`
	`beacon-inc-wait-time`
		wlsxAPImpersonation	126006, 127006

Detecting AP Spoofing on page 602	`ids impersonation-profile <profile-name>`
	`detect-ap-spoofing`
	`ap-spoofing-quiet-time`
		wlsxAPSpoofingDetected
		wlsxClientAssociatingOnWrongChannel
			126069, 126070, 127069, 127070

Feature	Command	Trap	Syslog ID
Detecting Bad WEP Initialization on page 602	`ids unauthorized-device-profile <profile-name> detect-bad-wep`	wlsxRepeatWEPIVViolation wlsxStaRepeatWEPIVViolation wlsxWeakWEPIVViolation wlsxStaWeakWEPIVViolation	126014, 126015, 126016, 126017, 127014, 127015, 127016, 127017
Detecting a Beacon Frame Spoofing Attack on page 602	`ids impersonation-profile <profile-name> detect-beacon-wrong-channel beacon-wrong-channel-quiet-time`	wlsxMalformedFrameWrongChannel Detected	126086, 127086
Detecting a Client Flood Attack on page 602	`ids dos-profile <profile-name> detect-client-flood client-flood-threshold client-flood-inc-time client-flood-quiet-time`	wlsxClientFloodAttack	126064, 127064
Detecting a CTS Rate Anomaly on page 602	`ids dos-profile <profile-name> detect-cts-rate-anomaly cts-rate-threshold cts-rate-time-interval cts-rate-quiet-time`	wlsxCtsRateAnomaly	126073, 127073
Detecting Devices with an Invalid MAC OUI on page 603	`ids unauthorized-device-profile <profile-name> detect-invalid-mac-oui mac-oui-quiet-time`	wlsxInvalidMacOUIAP wlsxInvalidMacOUISta	126029, 126030, 127029, 127030
Detecting an Invalid Address Combination on page 603	`ids dos-profile <profile-name> detect-invalid-address-combination invalid-address-combination-quiet-time`	wlsxInvalidAddressCombination	126079, 127079
Detecting an Overflow EAPOL Key on page 603	`ids dos-profile <profile-name> detect-overflow-eapol-key overflow-eapol-key-quiet-time`	wlsxMalformedOverflowEAPOLKey Detected	126082, 127082
Detecting Overflow IE Tags on page 603	`ids dos-profile <profile-name> detect-overflow-ie overflow-ie-quiet-time`	wlsxOverflowIEDetected	126084, 127084
Detecting a Malformed Frame-Assoc Request on page 603	`ids dos-profile <profile-name> detect-malformed-assoc-req malformed-assoc-req-quiet-time`	wlsxMalformedAssocReqDetected	126080, 127080

Feature	Command	Trap	Syslog ID

Detecting Malformed Frame-Auth on page 603	`ids dos-profile <profile-name>`
	`detect-malformed-frame-auth`
	`malformed-auth-frame-quiet-time`
		wlsxMalformedAuthFrameDetected	126083, 127083

Detecting a Malformed Frame-HT IE on page 603	`ids dos-profile <profile-name>`
	`detect-malformed-htie`
	`malformed-htie-quiet-time`
		wlsxMalformedHTIEDetected	126081, 127081

Detecting a Malformed Frame-Large Duration on page 603	`ids-dos-profile <profile-name>`
	`detect-malformed-large-duration`
	`malformed-large-duration-quiet-time`
		wlsxMalformedFrameLargeDuration Detected	126085, 127085

Detecting a Misconfigured AP on page 603 (WEP, WPA, SSID, Channel, OUI)	`ids unauthorized-device-profile <profile-name>`
	`detect-misconfigured-ap`
	`privacy`
	`require-wpa`
	`valid-and-protected-ssid`
	`cfg-valid-11g-channel`
	`cfg-valid-11a-channel`
	`valid-oui`
		wlsxWEPMisconfiguration
		wlsxWPAMisconfiguration
		wlsxSSIDMisconfiguration
		wlsxChannelMisconfiguration
		wlsxOUIMisconfiguration
			126011, 126028, 126010, 126008, 126009, 127011, 127028, 127010, 127008, 127009

Detecting a CTS Rate Anomaly on page 602	`ids dos-profile <profile-name>`
	`detect-rts-rate-anomaly`
	`rts-rate-threshold`
	`rts-rate-time-interval`
	`rts-rate-quiet-time`
		wlsxRtsRateAnomaly	126074, 127074

Detecting a Windows Bridge on page 604	`ids unauthorized-device-profile <profile-name>` `detect-windows-bridge`	wlsxWindowsBridgeDetectedAP wlsxWindowsBridgeDetectedSta wlsxNAdhocNetworkBridgeDetected AP wlsxNAdhocNetworkBridgeDetected Sta	126039, 126040, 126041, 126042, 127039, 127040, 127041, 127042

Detecting a Wireless Bridge on page 604	`ids unauthorized-device-profile <profile-name>`
	`detect-wireless-bridge`
	`wireless-bridge-quiet-time`
		wlsxWirelessBridge	126036, 127036

Detecting Broadcast Deauthentication on page 604	`ids signature-matching-profile <profile-name>`
	`signature deauth-Broadcast`
	
	`ids general-profile <profile-name>`
	`signature-quiet-time`
		wlsxNSignatureMatchDeauthBcast	126047, 127047

Feature	Command	Trap	Syslog ID

Detecting Broadcast Disassociation on page 604	`ids signature-matching-profile <profile-name>`
	`signature disassoc-Broadcast`
	
	`ids general-profile <profile-name>`
	`signature-quiet-time`
		wlsxNSignatureMatchDisassocBcast	126066, 127066

Detecting Netstumbler on page 604	`ids signature-matching-profile <profile-name>`
	`signature 'Netstumbler Generic'`
	`signature 'Netstumbler Version 3.3.0.x'`
	
	`ids general-profile <profile-name>`
	`signature-quiet-time`
		wlsxNSignatureMatchNetstumbler	126043, 127043

Detecting Valid SSID Misuse on page 604	`ids-unauthorized-device-profile <profile-name>`
	`detect-valid-ssid-misuse`
	`valid-and-protected-ssid`
		wlsxValidSSIDViolation	126007, 127007

Detecting Wellenreiter on page 604	`ids signature-matching-profile <profile-name>`
	`signature Wellenreiter`
	
	`ids general-profile <profile-name>`
	`signature-quiet-time`
		wlsxNSignatureMatchWellenreiter	126067, 127067

Detecting an 802.11n 40 MHz Intolerance Setting

When a client sets the HT capability intolerant bit to indicate that it is unable to participate in a 40 MHz BSS, the AP must use lower data rates with all of its clients. Network administrators often want to know if there are devices that are advertising 40 MHz intolerance, as this can impact the performance of the network.

Detecting Active 802.11n Greenfield Mode

When 802.11 devices use the HT operating mode, they can not share the same channel as 802.11a/b/g stations. Not only can they not communicate with legacy devices, the way they use the transmission medium is different, which would cause collisions, errors, and retransmissions.

Detecting Ad hoc Networks

An ad hoc network is a collection of wireless clients that form a network amongst themselves without the use of an AP. As far as network administrators are concerned, ad hoc wireless networks are uncontrolled. If they do not use encryption, they may expose sensitive data to outside eavesdroppers. If a device is connected to a wired network and has bridging enabled, an ad hoc network may also function like a rogue AP. Additionally, ad hoc networks can expose client devices to viruses and other security vulnerabilities. For these reasons, many administrators choose to prohibit ad hoc networks.

Detecting an Ad hoc Network Using a Valid SSID

If an unauthorized ad hoc network is using the same SSID as an authorized network, a valid client may be tricked into connecting to the wrong network. If a client connects to a malicious ad hoc network, security breaches or attacks can occur.

Detecting an AP Flood Attack

Fake AP is a tool that was originally created to thwart wardrivers by flooding beacon frames containing hundreds of different addresses. This would appear to a wardriver as though there were hundreds of APs in the area, thus concealing the real AP. An attacker can use this tool to flood an enterprise or public hotspots with fake AP beacons to confuse legitimate users and to increase the amount of processing need on client operating systems.

Detecting AP Impersonation

In AP impersonation attacks, the attacker sets up an AP that assumes the BSSID and ESSID of a valid AP or a neighboring AP. AP impersonation attacks can be done for man-in-the-middle attacks, a rogue AP attempting to bypass detection, or a honeypot attack.

Detecting AP Spoofing

An AP Spoofing attack involves an intruder sending forged frames that are made to look like they are from a legitimate AP. It is trivial for an attacker to do this, since tools are readily available to inject wireless frames with any MAC address that the user desires. Spoofing frames from a legitimate AP is the foundation of many wireless attacks.

Detecting Bad WEP Initialization

This is the detection of WEP initialization vectors that are known to be weak. A primary means of cracking WEP keys is to capture 802.11 frames over an extended period of time and searching for such weak implementations that are still used by many legacy devices.

Detecting a Beacon Frame Spoofing Attack

In this type of attack, an intruder spoofs a beacon packet on a channel that is different from that advertised in the beacon frame of the AP. The Beacon wrong channel event is supported only on the 5 GHz band.

Detecting a Client Flood Attack

There are fake AP tools that can be used to attack wireless intrusion detection itself by generating a large number of fake clients that fill internal tables with fake information. If successful, it overwhelms the wireless intrusion system, resulting in a DoS.

Detecting a CTS Rate Anomaly

The RF medium can be reserved via Virtual Carrier Sensing using a CTS transaction. The transmitter station sends a RTS frame to the receiver station. The receiver station responds with a CTS frame. All other stations that receive these CTS frames will refrain from transmitting over the wireless medium for an amount of time specified in the duration fields of these frames.

Attackers can exploit the Virtual Carrier Sensing mechanism to launch a DoS attack on the WLAN by transmitting numerous RTS and/or CTS frames. This causes other stations in the WLAN to defer transmission to the wireless medium. The attacker can essentially block the authorized stations in the WLAN with this attack.

Detecting an RTS Rate Anomaly

The RF medium can be reserved via Virtual Carrier Sensing using an RTS transaction. The transmitter station sends a RTS frame to the receiver station. The receiver station responds with a CTS frame. All other stations that receive these RTS frames will refrain from transmitting over the wireless medium for an amount of time specified in the duration fields of these frames.

Attackers can exploit the Virtual Carrier Sensing mechanism to launch a DoS attack on the WLAN by transmitting numerous RTS and/or CTS frames. This causes other stations in the WLAN to defer

transmission to the wireless medium. The attacker can essentially block the authorized stations in the WLAN with this attack.

Detecting Devices with an Invalid MAC OUI

The first three bytes of a MAC address, known as the MAC OUI, is assigned by the IEEE to known manufacturers. Often, clients using a spoofed MAC address do not use a valid OUI and instead use a randomly generated MAC address.

Detecting an Invalid Address Combination

In this attack, an intruder can cause an AP to transmit deauthentication and disassociation frames to all of its clients. Triggers that can cause this condition include the use of broadcast or multicast MAC address in the source address field.

Detecting an Overflow EAPOL Key

Some wireless drivers used in access points do not correctly validate the EAPOL key fields. A malicious EAPOL-Key packet with an invalid advertised length can trigger a DoS or possible code execution. This can only be achieved after a successful 802.11 association exchange.

Detecting Overflow IE Tags

Some wireless drivers used in access points do not correctly parse the vendor-specific IE tags. A malicious association request sent to the AP containing an IE with an inappropriate length (too long) can cause a DoS and potentially lead to code execution. The association request must be sent after a successful 802.11 authentication exchange.

Detecting a Malformed Frame-Assoc Request

Some wireless drivers used in access points do not correctly parse the SSID information element tag contained in association request frames. A malicious association request with a null SSID (that is, zero length SSID) can trigger a DoS or potential code execution condition on the targeted device.

Detecting Malformed Frame-Auth

Malformed 802.11 authentication frames that do not conform to the specification can expose vulnerabilities in some drivers that have not implemented proper error checking. This feature checks for unexpected values in an Authentication frame.

Detecting a Malformed Frame-HT IE

The IEEE 802.11n HT IE is used to convey information about the 802.11n network. An 802.11 management frame containing a malformed HT IE can crash some client implementations, potentially representing an exploitable condition when transmitted by a malicious attacker.

Detecting a Malformed Frame-Large Duration

The virtual carrier-sense attack is implemented by modifying the 802.11 MAC layer implementation to allow random duration values to be sent periodically. This attack can be carried out on the ACK, data, RTS, and CTS frame types by using large duration values. This attack can prevent channel access to legitimate users.

Detecting a Misconfigured AP

A list of parameters can be configured to define the characteristics of a valid AP. This feature is primarily used when non-Aruba APs are used in the network, since the Aruba devices cannot configure the third-party APs. These parameters include WEP, WPA, OUI of valid MAC addresses, valid channels, and valid SSIDs.

Detecting a Windows Bridge

A Windows Bridge occurs when a client that is associated to an AP is also connected to the wired network, and has enabled bridging between these two interfaces.

Detecting a Wireless Bridge

Wireless bridges are normally used to connect multiple buildings together. However, an attacker could place (or have an authorized person place) a wireless bridge inside the network that would extend the corporate network somewhere outside the building. Wireless bridges are somewhat different from rogue APs, in that they do not use beacons and have no concept of association. Most networks do not use bridges – in these networks, the presence of a bridge is a signal that a security problem exists.

Detecting Broadcast Deauthentication

A deauthentication broadcast attempts to disconnect all stations in range. Rather than sending a spoofed deauth to a specific MAC address, this attack sends the frame to a broadcast address.

Detecting Broadcast Disassociation

By sending disassociation frames to the broadcast address (FF:FF:FF:FF:FF:FF), an attacker can disconnect all stations on a network for a widespread DoS.

Detecting Netstumbler

NetStumbler is a popular wardriving application used to locate 802.11 networks. When used with certain NICs, NetStumbler generates a characteristic frame that can be detected. Version 3.3.0 of NetStumbler changed the characteristic frame slightly.

Detecting Valid SSID Misuse

If an unauthorized AP (neighbor or interfering) is using the same SSID as an authorized network, a valid client may be tricked into connecting to the wrong network. If a client connects to a malicious network, security breaches or attacks can occur.

Detecting Wellenreiter

Wellenreiter is a passive wireless network discovery tool used to compile a list of APs along with their MAC address, SSID, channel, and security setting in the vicinity. It passively sniffs wireless traffic, and with certain version (versions 1.4, 1.5, and 1.6), sends active probes that target known default SSIDs.

Understanding Client Intrusion Detection

Generally, clients are more vulnerable to attacks than APs. Clients are more apt to associate with a malignant AP due to the client's driver behavior or a misconfigured client. It is important to monitor authorized clients to track their associations and to track any attacks raised against the client. Client attack detection is categorized as:

- **Detecting attacks against Aruba APs clients:** An attacker can perform an active DOS attack against an associated client, or perform a replay attack to obtain the keys of transmission which could lead to more serious attacks.
- **Monitoring Authorized clients:** Since clients are easily tricked into associating with unauthorized APs, tracking all misassociations of authorized clients is very important.

An authorized client is a client authorized to use the WLAN network. In ArubaOS, an authorized client is called a valid-client. ArubaOS automatically learns a valid client. A client is determined to be valid if it is associated to an authorized or valid AP using encryption; either Layer 2 or IPsec.



Detection of attacks is limited to valid clients and clients associated to valid APs. Clients that are associated as guests using unencrypted association are included in the attack detection. However, clients on neighboring (interfering) APs are not tracked for attack detection unless they are specified as valid.

Table 109 presents a summary of the client intrusion detection features with their related commands, traps, and syslog identification. Details of each feature follow the table.

Table 109: *Client Detection Summary*

Feature	Command	Trap	Syslog ID

Detecting a Block ACK DoS on page 606	`ids-dos-profile <profile-name>`
	`detect-block-ack-attack`
	`block-ack-quiet-time`
		wlsxBlockAckAttackDetected	126087, 127087

Detecting a ChopChop Attack on page 606	`ids-dos-profile <profile-name>`
	`detect-chopchop-attack`
	`chopchop-quiet-time`
		wlsxChopChopAttackDetected	126078, 127078

Detecting a Disconnect Station Attack on page 607	`ids dos-profile <name>`
	`detect-disconnect-sta`
	`disconnect-sta-quiet-time`
	`disconnect-sta-assoc-resp-threshold`
	`disconnect-deauth-disassoc-threshold`
		wlsxNDisconnectStationAttack	126035, 127035

Detecting an EAP Rate Anomaly on page 607	`ids-dos-profile <profile-name>`
	`detect-eap-rate-anomaly`
	`eap-rate-threshold`
	`eap-rate-time-interval`
	`eap-rate-quiet-time`
		wlsxEAPRateAnomaly	126032, 127032

Detecting a FATA-Jack Attack Structure on page 607	`ids dos-profile <profile-name>`
	`detect-fatajack-attack`
	`fatajack-attack-quiet-time`
		wlsxFataJackAttackDetected	126072, 127072

Detecting a Hotspotter Attack on page 607	`ids impersonation-profile <profile-name>`
	`detect-hotspotter-attack`
	`hotspotter-quiet-time`
		wlsxHotspotterAttackDetected	126088, 127088

Detecting a Meiners Power Save DoS Attack on page 607	`ids dos-profile <profile-name>`
	`detect-power-save-dos-attack`
	`power-save-dos-min-frames`
	`power-save-dos-quiet-time`
	`power-save-dos-threshold`
		wlsxPowerSaveDoSAttack	126109, 127109

Detecting an Omerta Attack on page 607	`ids dos-profile <profile-name>`
	`detect-omerta-attack`
	`omerta-attack-threshold`
	`omerta-attack-quiet-time`
		wlsxOmertaAttack	126071, 127071

Feature	Command	Trap	Syslog ID

Detecting Rate Anomalies on page 607	`ids dos-profile <profile-name> detect-rate-anomalies`
	
	`assoc-rate-thresholds`
	`disassoc-rate-thresholds`
	`deauth-rate-thresholds`
	`probe-request-rate-thresholds`
	`probe-response-rate-thresholds`
	`auth-rate-thresholds`
		wlsxChannelRateAnomaly
		wlsxNodeRateAnomalyAP
		wlsxNodeRateAnomalySta
			126061, 126062, 126063, 127061, 127062, 127063

Detecting a TKIP Replay Attack on page 607	`ids dos-profile detect-tkip-replay-attack tkip-replay-quiet-time`	wlsxTkipReplayAttackDetected	126077, 127077
Detecting Unencrypted Valid Clients on page 608	`ids unauthorized-device-profile detect-unencrypted-valid-client unencrypted-valid-client-quiet-time`	wlsxValidClientNotUsingEncryption	126065, 127065
Detecting a Valid Client Misassociation on page 608	`ids unauthorized-device-profile detect-valid-client-misassociation`	wlsxValidClientMisassociation	126075, 127075

Detecting an AirJack Attack on page 608	`ids signature-matching-profile signature AirJack`
	
	`ids general-profile signature-quiet-time`
		wlsxNSignatureMatchAirjack	126046, 127046

Detecting ASLEAP on page 608	`ids signature-matching-profile signature ASLEAP`
	
	`ids general-profile signature-quiet-time`
		wlsxNSignatureMatchAsleap	126044, 127044

Detecting a Null Probe Response on page 608	`ids signature-matching-profile signature Null Probe Response`
	
	`ids general-profile signature-quiet-time`
		wlsxNSignatureMatchNullProbeResp	126045, 127045

Detecting a Block ACK DoS

The Block ACK mechanism that was introduced in 802.11e, and enhanced in 802.11nD3.0, has a built-in DoS vulnerability. The Block ACK mechanism allows for a sender to use the ADDBA request frame to specify the sequence number window that the receiver should expect. The receiver will only accept frames in this window.

An attacker can spoof the ADDBA request frame causing the receiver to reset its sequence number window and thereby drop frames that do not fall in that range.

Detecting a ChopChop Attack

ChopChop is a plaintext recovery attack against WEP encrypted networks. It works by forcing the plaintext, one byte at a time, by truncating a captured frame and then trying all 256 possible values for

the last byte with a corrected CRC. The correct guess causes the AP to retransmit the frame. When that happens, the frame is truncated again.

Detecting a Disconnect Station Attack

A disconnect attack can be launched in many ways; the end result is that the client is effectively and repeatedly disconnected from the AP.

Detecting an EAP Rate Anomaly

To authenticate wireless clients, WLANs may use 802.1X, which is based on a framework called EAP. After an EAP packet exchange, and the user is successfully authenticated, the EAP-Success is sent from the AP to the client. If the user fails to authenticate, an EAP-Failure is sent. In this attack, EAP-Failure or EAP-Success frames are spoofed from the access point to the client to disrupting the authentication state on the client. This confuses the client's state, causing it to drop the AP connection. By continuously sending EAP Success or Failure messages, an attacker can effectively prevent the client from authenticating with the APs in the WLAN.

Detecting a FATA-Jack Attack Structure

FATA-Jack is an 802.11 client DoS tool that tries to disconnect targeted stations using spoofed authentication frames that contain an invalid authentication algorithm number.

Detecting a Hotspotter Attack

The Hotspotter attack is an evil-twin attack which attempts to lure a client to a malicious AP. Many enterprise employees use their laptop in Wi-Fi area hotspots at airports, cafes, malls etc. They have SSIDs of their hotspot service providers configured on their laptops. The SSIDs used by different hotspot service providers are well known. This enables the attackers to set up APs with hotspot SSIDs in close proximity of the enterprise premises. When the enterprise laptop Client probes for hotspot SSIDs, these malicious APs respond and invite the client to connect to them. When the client connects to a malicious AP, a number of security attacks can be launched on the client. Airsnarf is a popular hacking tool used to launch these attacks.

Detecting a Meiners Power Save DoS Attack

To save on power, wireless clients will sleep periodically, during which they cannot transmit or receive. A client indicates its intention to sleep by sending frames to the AP with the Power Management bit ON. The AP then begins buffering traffic bound for that client until it indicates that it is awake. An intruder could exploit this mechanism by sending (spoofed) frames to the AP on behalf of the client to trick the AP into believing the client is asleep. This will cause the AP to buffer most, if not all, frames destined for the client.

Detecting an Omerta Attack

Omerta is an 802.11 DoS tool that sends disassociation frames to all stations on a channel in response to data frames. The Omerta attack is characterized by disassociation frames with a reason code of 0x01. This reason code is unspecified and is not used under normal circumstances.

Detecting Rate Anomalies

Many DoS attacks flood an AP or multiple APs with 802.11 management frames. These can include authenticate or associate frames, which are designed to fill up the association table of an AP. Other management frame floods, such as probe request floods, can consume excess processing power on the AP.

Detecting a TKIP Replay Attack

TKIP is vulnerable to replay (via WMM or QoS) and plain text discovery (via ChopChop). This affects all WPA-TKIP usage. By replaying a captured TKIP data frame on other QoS queues, an attacker can manipulate the RC4 data and checksum to derive the plain text at a rate of one byte per minute.

By targeting an ARP frame and guessing the known payload, an attacker can extract the complete plain text and MIC checksum. With the extracted MIC checksum, an attacker can reverse the MIC AP to Station key and sign future messages as MIC compliant, opening the door for more advanced attacks.

Detecting Unencrypted Valid Clients

An authorized (valid) client that is passing traffic in unencrypted mode is a security risk. An intruder can sniff unencrypted traffic (also known as packet capture) with software tools known as sniffers. These packets are then reassembled to produce the original message.

Detecting a Valid Client Misassociation

This feature does not detect attacks, but rather it monitors authorized (valid) wireless clients and their association within the network. Valid client misassociation is potentially dangerous to network security. The four types of misassociation that we monitor are:

- **Authorized Client associated to Rogue:** A valid client that is associated to a rogue AP.
- **Authorized Client associated to External AP:** An external AP, in this context, is any AP that is not valid and not a rogue.
- **Authorized Client associated to Honeypot AP:** A honeypot is an AP that is not valid but is using an SSID that has been designated as valid or protected.
- **Authorized Client in ad-hoc connection mode:** A valid client that has joined an ad-hoc network.

Detecting an AirJack Attack

AirJack is a suite of device drivers for 802.11(a/b/g) raw frame injection and reception. It was intended to be used as a development tool for all 802.11 applications that need to access the raw protocol. However, one of the tools included allowing users to force all users off an AP.

Detecting ASLEAP

ASLEAP is a tool created for Linux systems used to attack Cisco LEAP authentication protocol.

Detecting a Null Probe Response

A null probe response attack has the potential to crash or lock up the firmware of many 802.11 NICs. In this attack, a client probe-request frame will be answered by a probe response containing a null SSID. A number of popular NIC cards will lock up upon receiving such a probe response.

Configuring Intrusion Protection

Intrusion protection features support containment of an AP or a client. In the case of an AP, we will attempt to disconnect all clients that are connected or attempting to connect to the AP. In the case of a client, the client's association to an AP is targeted. The following containment mechanisms are supported:

- **Deauthentication containment:** An AP or client is contained by disrupting its association on the wireless interface.
- **Tarpit containment:** An AP is contained by luring clients that are attempting to associate with it to a tarpit. The tarpit can be on the same channel as the AP being contained, or on a different channel

(see Tarpit Shielding Overview on page 625).

- **Wired containment:** An AP or client is contained by disrupting its connection on the wired interface.

The WIP feature supports separate enforcement policies that use the underlying containment mechanisms to contain an AP or a client that do not conform to the policy. These policies are discussed in the sections that follow:

- Understanding Infrastructure Intrusion Protection on page 609
- Understanding Client Intrusion Protection on page 611
- Warning Message for Containment Features on page 611

Understanding Infrastructure Intrusion Protection

The following is the list of infrastructure intrusion protection features with their related commands:

- Protecting 40 MHz 802.11 High Throughput Devices
- Protecting 802.11n High Throughput Devices
- Protecting Against Adhoc Networks
- Protecting Against Adhoc Networks Using Valid SSID
- Protecting Against AP Impersonation
- Protecting Against Misconfigured APs
- Protecting Against Wireless Hosted Networks
- Protecting SSIDs
- Protecting Against Rogue Containment
- Protecting Against Suspected Rogue Containment
- Protection Against Wired Rogue APs

Protecting 40 MHz 802.11 High Throughput Devices

Protection from AP(s) that support 40 MHz HT involves containing the AP such that clients can not connect.

To enable this feature, enter the following command in the configuration mode of the CLI:

```
ids unauthorized-device-profile <profile-name>
protect-ht-40mhz
```

Protecting 802.11n High Throughput Devices

Protection from APs that support HT involves containing the AP such that the clients can not connect.

To enable this feature, enter the following command in the configuration mode of the CLI:

```
ids unauthorized-device-profile <profile-name>
protect-high-throughput
```

Protecting Against Adhoc Networks

Protection from an adhoc network involves containing the adhoc network so that clients cannot connect to it. The basic adhoc protection feature protects against adhoc networks using WPA or WPA2 security. The enhanced adhoc network protection feature protects against open or WEP adhoc networks. Both features can be used together for maximum protection, or enabled or disabled separately.

To enable this feature, enter the following command in the configuration mode of the CLI:

```
ids unauthorized-device-profile <profile-name>
protect-adhoc-enhanced
protect-adhoc-network
```



This feature requires that you enable the **wireless-containment** setting in the IDS general profile.

Protecting Against Adhoc Networks Using Valid SSID

Protection from adhoc networks using valid SSID involves containing the adhoc networks that use a valid or protected SSIDs so that clients cannot connect to it. This feature provides protection against WPA, WPA2, WEP, or open adhoc networks.

To enable this feature, enter the following command in the configuration mode of the CLI:

```
ids unauthorized-device-profile <profile-name>
protect-adhoc-using-valid-ssid
```

Protecting Against AP Impersonation

Protection from AP impersonation involves containing both the legitimate and impersonating AP so that clients cannot connect to either AP.

To enable this feature, enter the following command in the configuration mode of the CLI:

```
ids impersonation-profile <profile-name>
protect-ap-impersonation
```

Protecting Against Misconfigured APs

Protect Misconfigured AP enforces that valid APs are configured properly. An offending AP is contained by preventing clients from associating to it.

To enable this feature, enter the following command in the configuration mode of the CLI:

```
ids unauthorized-device-profile <profile-name>
protect-misconfigured-ap
```

Protecting Against Wireless Hosted Networks

Clients using the Windows wireless hosted network feature can act as an access point to which other wireless clients can connect, effectively becoming a Wi-Fi HotSpot. This creates a security issue for enterprises, because unauthorized users can use a hosted network to gain access to the corporate network, and valid users that connect to a hosted network are vulnerable to attacks or security breaches. This feature detects a wireless hosted network, and contains the client hosting this network.

To enable this feature, enter the following command in the configuration mode of the CLI:

```
ids unauthorized-device-profile <profile-name>
detect-wireless-hosted-network protect-wireless-hosted-network
```

Protecting SSIDs

Protect SSID enforces that valid or protected SSIDs are used only by valid APs. An offending AP is contained by preventing clients from associating to it.

To enable this feature, enter the following command in the configuration mode of the CLI:

```
ids unauthorized-device-profile <profile-name>
protect-ssid
```

Protecting Against Rogue Containment

By default, rogue APs are not automatically disabled. Rogue containment automatically disables a rogue AP by preventing clients from associating to it.

To enable this feature, enter the following command in the configuration mode of the CLI:

```
ids unauthorized-device-profile <profile-name>
rogue-containment
```

Protecting Against Suspected Rogue Containment

By default, suspected rogue APs are not automatically contained. In combination with the suspected rogue containment confidence level, suspected rogue containment automatically disables a suspect rogue by preventing clients from associating to it.

To enable this feature, enter the following command in the configuration mode of the CLI:

```
ids unauthorized-device-profile
suspect-rogue-containment
suspect-rogue-conf-level
```

Protection Against Wired Rogue APs

This feature enables containment from the wired side of the network. The basic wired containment feature in the IDS general profile isolates layer-3 APs whose wired interface MAC addresses are the same as (or one character off from) their BSSIDs. The enhanced wired containment feature can also identify and contain an AP with a preset wired MAC address that is completely different from the AP's BSSID. In many non-Aruba APs, the MAC address the AP provides to wireless clients as a 'gateway MAC' is offset by one character from its wired MAC address. This enhanced feature allows ArubaOS to check to see if a suspected Layer-3 rogue AP's MAC address follows this common pattern.

To enable this feature, enter the following command in the configuration mode of the CLI:

```
ids general-profile
wired-containment
wired-containment-ap-adj-mac
wired-containment-susp-l3-rogue
```

Understanding Client Intrusion Protection

The following are the client intrusion protection features with their related commands.

Protecting Valid Stations

Protecting a valid client involves disconnecting that client if it is associated to a non-valid AP.

To protect valid stations access the CLI and enter the following command in the configuration mode:

```
ids unauthorized-device-profile <profile-name> protect-valid-sta
```

Protecting Windows Bridge

Protecting from a Windows Bridge involves containing the client that is forming the bridge so that it can not connect to the AP.

To protect Windows bridge access the CLI and enter the following command in the configuration mode:

```
ids unauthorized-device-profile <profile-name> protect-windows-bridge
```

Warning Message for Containment Features

The feature for enabling wireless containment under the **IDS Unauthorized Device** profile and **IDS Impersonation** profile may be in violation of certain FCC regulatory statutes. To address this, a warning message is issued each time the command is enabled:

- If enabled through the WebUI, the warning message will appear before the command is executed.
- If enabled through the CLI, the warning message will appear after the command is executed.

Configuring the WLAN Management System

The WLAN management system on Mobility Master monitors wireless traffic to detect any new AP or wireless client station in the RF environment. When an AP or wireless client is detected, it is classified, and its classification is used to determine the security policies that should be enforced on the AP or

client. By default, the WMS service is terminated at Mobility Master, which requires every AP across the network to communicate with the WMS service on Mobility Master.

ArubaOS now allows IPv6 support for WMS and Air Monitor features of WIPS in both dual-stack and native IPv6 deployments.

This section describes the following topics:

- Configuring General WMS settings
- Configuring Local WMS Settings
- Mobility Master WMS Termination vs. Managed Device WMS Termination
- Managing the WMS Database

Configuring General WMS settings

Use the IDS WMS General profile to configure general WMS settings such as AP ageout times and update intervals, and enable the collection of statistics for monitored APs and clients.

The following procedure describes how to configure general WMS settings:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. Expand the **IDS** menu and select **IDS WMS General**.
3. Configure the parameters as described in Table 110 and then click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Table 110: *WMS Configuration Parameters*

Parameter	Description
AP poll interval	Interval, in milliseconds, for communication between the managed device and Aruba APs. The managed device contacts the AP at this interval to download AP to station associations, update policy configuration changes, and download AP and station statistics. Default: 60000 milliseconds (1 minute)
AP poll retries	Maximum number of failed polling attempts before the polled AP is considered to be down. Default: 3
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
Adhoc AP ageout interval	The amount of time, in minutes, that an ad hoc (IBSS) AP unseen by any problems before it is deleted from the database. Enter 0 to disable ageout. Default: 5 minutes
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
Statistics update	Enables or disables statistics update in the database. Default: disabled

Parameter	Description
Persistent Neighbor APs	When enabled, this feature prevents APs that are marked as neighbor APs from being aged out. Default: disabled
Persistent Valid STAs	When enabled, this feature prevents valid stations from being aged out. Default: disabled
AP learning	Enables or disables AP learning. Learning affects the way APs are classified. Default: disabled
Propagate Wired Macs	Enables the propagation of the gateway wired MAC information. Default: enabled
Collect Stats for Monitored APs and Clients	Enables collection of statistics (up to 25,000 entries) on Mobility Master for monitored APs and clients. Default: disabled
Learn System Wired Macs:	Enable or disable “learning” of wired MACs on the managed device. Default: disabled

The following CLI command configures WMS. The parameters in this command are described in detail in Table 110.

```
(host)[mynode]config# ids wms-general-profile
```

Configuring Local WMS Settings

The configuration parameters in IDS WMS local system profile allow the user to change the default behavior and table sizes of the WMS on specific managed devices.

The following procedure describes how to configure local WMS settings:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. Expand the **IDS** menu and select **IDS WMS Local System**.
3. Configure the parameters as described in Table 111 and then click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Table 111: *IDS WMS Local System Profile Settings*

Parameter	Description
Max AP Threshold	Set the max threshold for the total number of APs
Max STA Threshold	Set the max threshold for the total number of stations.
Max RBTree Entries	Set the max threshold for the total number of AP and station RBTree entries.
Max System Wired MACs	Set the max number of system wired MAC table entries learned by the managed device.

Parameter	Description
Override Service Termination	Override the system-determined termination mode, and terminate WMS service at the managed device to which the AP is associated. Do not use this option if you have multiple managed devices in one location, as WMS will not operate correctly. For more information, see Configuring Local WMS Settings on page 613.
Periodic AP Snapshot Interval	Set the interval in minutes at which to generate a periodic snapshot of monitored APs. The (AMON) messages comprising the snapshot will be spread over this interval.
Periodic Rogue AP Snapshot Interval	Set the interval in minutes at which to generate a periodic snapshot of monitored Rogue APs. The (AMON) messages comprising the snapshot will be spread over this interval.
Periodic STA Snapshot Interval	Set the interval in minutes at which to generate a periodic snapshot of monitored clients. The (AMON) messages comprising the snapshot will be spread over this interval.
System Wired MAC Update Interval	Set the interval, in minutes, for repopulating the system wired MAC table at the managed device.

The following CLI command configures local WMS settings**.** The parameters in this command are described in detail in Table 111.

```
(host) [mynode] (config) #ids wms-local-system-profile
```

Mobility Master WMS Termination vs. Managed Device WMS Termination

By default, the devices and events detected by a managed device are sent to Mobility Master, allowing Mobility Master to update its database with AP, client, and event information from that managed device. This is the recommended mode for terminating WMS services.

If a managed device is installed at a location with strict bandwidth limitations, the WMS services can optionally be configured to terminate at the managed device. Local managed device termination of WMS services must be enabled with caution. Enabling this feature reduces the bandwidth used by messages between the managed device and Mobility Master, but does introduce some serious limitations. Optimal device classification and IDS detection or protection requires a centralized network-wide view that is best provided by WMS termination on Mobility Master.



Enable local (managed device) termination of the WMS service with caution, as enabling this feature may impact WMS device classification and IDS detection and protection on your network. **This feature is only supported on a network topology where the managed device is geographically away from another managed device terminating APs.**

This section describes the following topics:

- AMON Messaging between WMS on a Managed Device and WMS on Mobility Master on page 614
- Supported AMON Message Types on page 615

AMON Messaging between WMS on a Managed Device and WMS on Mobility Master

If you enable local termination of the WMS service on a managed device, the default behavior for this feature prevents Mobility Master from monitoring devices seen by APs on that managed device. As a

result, network administrators must view these devices via the managed device WebUI and command-line interfaces. This is the recommended method for monitoring devices at locations where a managed device locally terminates the WMS service.

However, a managed device can optionally be configured to send AMON messages with information about monitored devices and events to the WMS service on Mobility Master. The managed device does not send AMON messages to Mobility Master unless AMON messaging is enabled via Mobility Master WebUI.



Do not enable WMS AMON messages from a managed device to Mobility Master if the WMS service does not terminate on the managed device, as this will disrupt WMS functionality for the APs associated to that managed device. AMON messages sent from a managed device to Mobility Master are likely to consume a substantial amount of bandwidth, potentially eliminating and bandwidth savings provided by local termination of WMS on a managed device.

The following procedure describes how to allow a managed device that locally terminates the WMS service to send AMON messages to Mobility Master:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > More > General** page and configure the Mobility Master as management server for the managed device.
2. Enable AMON messaging from the managed device to Mobility Master using the **Mgmt Config** profile in **Controller Profile** under the **Configuration > System > Profiles** page of the Mobility Master WebUI.

Supported AMON Message Types

A managed device terminating the WMS service can send the following AMON message types to Mobility Master:

- **Monitored AP Info Messages**: This message is sent when a monitored AP is newly added to the WMS, or a monitored AP's classification, confidence level, SSID, or encryption type has changed. To enable this message type, select the **Monitored Info - Add/Update** check box in the **Mgmt Config** profile.
- **Monitored AP Delete Messages**: This message is not used by WMS. Although this message type can be enabled via the **Monitored Info - Deletion** option in the **Mgmt Config** profile, best practices is to keep this option disabled.
- **Monitored Station Info Message**s: This message is sent when a monitored client is newly added to the WMS, or a monitored client's BSSID or rogue station type has changed. To enable this message type, select the **Monitored Info - Add/Update** check box in the Mgmt Config profile.
- **Monitored Station Delete Messages**: This message is not used by WMS. Although this message type can be enabled via the **Monitored Info - Deletion** option in the **Mgmt Config** profile, best practices is to keep this option disabled.
- **Rogue AP Info Messages**: This message is sent when an AP is newly classified as a rogue or suspected Rogue, or when the AP confidence level changes. To enable this message type, select the **Monitored Info - Add/Update** check box in the Mgmt Config profile.
- **Wireless IDS Event Info Message**: (new): This message type sends information about Intrusion Detection System events as they are seen. To enable this message type, select the **Wireless IDS Event Info** check box in the **Mgmt Config** profile.
- **Periodic AP Snapshots**: This message type sends a snapshot of all monitored APs in WMS every update period. To enable this message type, select the **Monitored Info - Periodic Snapshot** check

box in the **Mgmt Config** profile. Use the **Periodic STA Snapshot Interval** parameter in the IDS WMS Local System profile to set the interval during which the AP snapshot messages are sent.

- **Periodic Station Snapshots**: This message type sends a snapshot of all monitored clients in WMS every update period. To enable this message type, select the **Monitored Info - Periodic Snapshot** check box in the **Mgmt Config** profile. Use the **Periodic Rogue AP Snapshot Interval** parameter in the IDS WMS Local System profile to set the interval during which the client snapshot messages are sent.
- **Periodic Rogue AP Snapshots**: This message type sends a snapshot of all rogue APs in WMS every update period. To enable this message type, select the **Monitored Info - Periodic Snapshot** check box in the **Mgmt Config** profile. Use the **Periodic AP Snapshot Interval** parameter in the IDS WMS Local System profile to set the interval during which the rogue AP snapshot messages are sent.

Managing the WMS Database

The WMS process interacts with all the AM processes in the network. When WMS receives an event message from an AM, the WMS process will save the event information along with the BSSID of the AP that generated the event in the WMS database. Use the following commands in the CLI to manage the WMS database.

The **wms export-db** command exports the specified file as an ASCII text file into the WMS database.

```
(host)[mynode]#wms export-db <filename>
```

The **wms import-db** command imports the specified file into the WMS database:

```
(host)[mynode]#wms import-db <filename>
```

The **wms reint-db** command reinitializes the WMS database. Note that this command does not make an automatic backup of the current database.

```
(host)[mynode] #wms reint-db
```

This section describes the following topics:

- Optimizing Classification Behavior on page 616
- Managing the WMS Database on page 616

Optimizing Classification Behavior

APs can be configured to periodically send WMS a list of monitored devices that are still unclassified. Once the WMS receives this list, a classification message is sent from the WMS to the AP, to classify each unclassified device.

The following procedure describes how to configure IDS General profile parameters:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles.**
2. Expand the **IDS** menu and select **IDS General.**
3. Expand the **Advanced** accordion.
4. Configure the parameters as described in Table 112 then click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Table 112: *IDS General Profile Parameters*

Parameter	Description
Unclassified AP Update	Enables or disables classification updates for monitored APs. If this option is enabled, it helps decrease the delay in the speed at which the devices are classified. Default: Disabled
Unclassified STA Update	Enables or disables classification updates for monitored clients. If this option is enabled, it helps decrease the delay in the speed at which the devices are classified. Default: Disabled

The following CLI commands configure IDS General Profile parameters:

```
(host) [md](config)# ids general-profile <profile-name>
(host) [md] (IDS General Profile "<profile-name>") #unclass-ap-update
(host) [md] (IDS General Profile "<profile-name>")unclass-device-update-interval
(host) [md] (IDS General Profile "<profile-name>")unclass-sta-update
```

Managing the List of Valid Exempt Clients

The network administrator can configure clients to be exempted from valid station protection and valid station misassociation detection by adding the mac-address of those devices to the valid-exempt-list.

Once a client MAC address is added to the valid-exempt list:

- If the client exists in the WMS, the classification is set to valid.
- If the client does not exist in the WMS, a client entry is created and then the classification is set to valid.
- After the classification is done, APs that are seeing the client are notified that the client is added to the valid-exempt list.



A maximum of 200 MAC addresses can be added to a valid-exempt list. The valid-exempt list is not retained after the managed device reboots or a process is restarted.

You can configure clients to be exempted from valid station protection and valid station misassociation using the CLI.

The following CLI commands add or remove MAC addresses from the valid-exempt list:

```
(host)[md](config) #wms client <macaddr> valid-exempt insert
(host)[md](config) #wms client <macaddr> valid-exempt remove
```

The following CLI command displays a list of configured valid-exempt clients:

```
(host)[md] #show wms client valid-exempt
```

The following CLI command displays a list of clients that are viewed by the AP and marked as valid-exempt:

```
(host)[md] #show ap monitor client-list ap-name <> valid-exempt
```

The following CLI command displays the number of MAC addresses added to the valid-exempt client list:

```
(host)[md] #show wms counters
  Counters
  --------
  Name                        Value
  ----                        -----
  DB Reads                    288268
  DB Writes                   350870
  Probe Table DB Reads        2477
```

```
Probe Table DB Writes       952
AP Table DB Reads           143992
AP Table DB Writes          138867
STA Table DB Reads          40404
STA Table DB Writes         99687
Probe STA Table DB Reads    101352
Probe STA Table DB Writes   117566
Probe Register              2476
Probe State Update          37077
Set RAP Type                42552
Set RAP Type Conf Level     152
```
Valid Exempt Station Macs 10

Understanding Client Blacklisting

When a client is blacklisted in the Aruba system, the client is not allowed to associate with any AP in the network for a specified amount of time. If a client is connected to the network when it is blacklisted, a deauthentication message is sent to force the client to disconnect. While blacklisted, the client cannot associate with another SSID in the network.

The managed device retains the client blacklist in the user database, so the information is not lost if the managed device reboots. When you import or export the managed device's user database, the client blacklist will be exported or imported as well.

ArubaOS now forwards the client blacklist to the database of all the managed devices from the Mobility Master, when the blacklist is managed through the WebUI. Hence, the configuration and monitoring of client blacklist is centralized at the Mobility Master in the WebUI.

This section describes the following topics:

- Methods of Blacklisting on page 618
- Setting Blacklist Duration on page 620
- Removing a Client from Blacklisting on page 621

Methods of Blacklisting

There are several ways in which a client can be blacklisted in the Aruba system:

- You can manually blacklist a specific client. See Blacklisting Manually on page 619 for more information.
- A client fails to successfully authenticate for a configured number of times for a specified authentication method. The client is automatically blacklisted. See Blacklisting by Authentication Failure on page 619 for more information.
- A DoS or man in the middle attack has been launched in the network. Detection of these attacks can cause the immediate blacklisting of a client. See Understanding Client Blacklisting on page 618 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 Mobility Master via an XML API server. When the managed device receives the client blacklist request from the server, it blacklists the client, logs an event, and sends an SNMP trap.

See External Services Interface on page 1195 for more information.



The External Services Interface feature requires the PEFNG license installed in the managed device.

Blacklisting Manually

There are several reasons why you may choose to blacklist a client. For example, you can enable different Aruba IDS features that detect suspicious activities, such as DoS attacks. When these activities are detected, an event is logged and an SNMP trap is sent with the client information. To blacklist a client, you need to know its MAC address.

ArubaOS now allows you to manage blacklisted clients in stand-alone controllers as well as in Mobility Masters and the following procedure describes how to manage blacklisted clients:

In the **Managed Network** node hierarchy, navigate to either the **Dashboard > Security** or **Dashboard > Overview** page:

1. (Optional) From the **Dashboard > Security** page:
   a. Click the **Blacklist** icon or donut chart area in the **Blacklist** window to open the **Blacklisted Clients** table.
   b. Select a client from the **Wireless Clients** table.

   The **Blacklist** icon and donut chart area remain inactive when there are no blacklisted clients available. However, you can click the **Detected Radios**, **Detected Clients**, or **Events** icon to open a new window that displays the **Blacklisted Clients** table.

   c. Click the **+** icon on the Action bar to open the **Add to Blacklist** pop-up window.
   d. In the **Add to Blacklist** pop-up window, enter the MAC address of the client, and click **Add**.
2. (Optional) From the **Dashboard > Overview** page:
   a. Click the **Clients** icon or donut chart area in the **Clients** window to open the **Wireless Clients** table.
   b. Select a client from the **Wireless Clients** table.
   c. Click the **+** icon on the Action bar to open the **Add to Blacklist** pop-up window.
   d. In the **Add to Blacklist** pop-up window, click **Add**.

The client is blacklisted and is listed in the **Blacklisted Clients** table.



When you manually blacklist a client from the Mobility Master, the client gets blacklisted permanently.

For more information about blacklisted clients, see Dashboard Monitoring on page 410.

The following CLI command manually blacklists a client:

```
(host) [md] #stm add-blacklist-client
```



The **stm add-blacklist-client** command configures the blacklist only on the local managed device and is not applicable to other managed devices.

Blacklisting by Authentication Failure

You can configure a maximum authentication failure threshold for each of the following authentication methods:

- 802.1X
- MAC

- Captive portal
- VPN

When a client exceeds the configured threshold for one of the above methods, the client is automatically blacklisted by the managed device, an event is logged, and an SNMP trap is sent. By default, the maximum authentication failure threshold is set to 0 for the above authentication methods, which means that there is no limit to the number of times a client can attempt to authenticate.

With 802.1X authentication, you can also configure blacklisting of clients who fail machine authentication.



When clients are blacklisted because they exceed the authentication failure threshold, they are blacklisted indefinitely by default. You can configure the duration of the blacklisting; see Setting Blacklist Duration on page 620.

The following procedure describes how to set the authentication failure threshold:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In **All Profiles** expand the **Wireless LAN** list, select the appropriate authentication profile, then select the profile instance.
3. Enter a value in the **Max Authentication failures** field.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands set the authentication failure threshold:

```
(host) [md] (config) #aaa authentication {captive-portal|dot1x|mac|vpn} <profile>
(host) [md] (<Auth-Profile> <profile-name>) # max-authentication-failures <number>
```

Setting Blacklist Duration

You can configure the duration that clients are blacklisted on a per-SSID basis via the virtual AP profile. There are two different blacklist duration settings:

- For clients that are blacklisted due to authentication failure. By default, this is set to 0 (the client is blacklisted indefinitely).
- For clients that are blacklisted due to other reasons, including manual blacklisting. By default, this is set to 3600 seconds (one hour). You can set this to 0 to blacklist clients indefinitely.

The following procedure describes how to configure the blacklist duration:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration >System > Profiles** page.
2. In **All Profiles**, select **Wireless LAN**, then **Virtual AP**. Select the virtual AP instance.
3. To set a blacklist duration for authentication failure, expand the **Advanced** accordion and enter a value for **Authentication Failure Blacklist Time**.
4. To set a blacklist duration for other reasons, expand the **Advanced** accordion and enter a value for **Blacklist Time**.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the blacklist duration:

```
(host) [md] (config) #wlan virtual-ap default
(host) [md] (Virtual AP profile "default") #auth-failure-blacklist-time <seconds>
(host) [md] (Virtual AP profile "default") #blacklist-time <seconds>
```

Removing a Client from Blacklisting

The following procedure describes how to manually remove one or multiple blacklisted clients from a managed device:

1. In the **Managed Network** node hierarchy, navigate to the **Dashboard > Security** page.
2. Click the **Blacklist** icon or donut chart area in the **Blacklist** window.
3. The **Blacklisted Clients** table is displayed.
4. Hover your mouse over the wireless client that you want to remove from the blacklist, and select the corresponding check box.
5. (Optional) Hover your mouse over multiple wireless clients that you want to remove from the blacklist, and select the corresponding check boxes.
6. Click the **Delete blacklisted client** icon.
7. The **Confirm Deletion** pop-up window is displayed.
8. Click **Delete** to delete the client(s) from the **Blacklisted Clients** table.

The following CLI command removes a client from blacklisting:

```
(host) [md] #stm remove-blacklist-client <macaddr>
```

The following CLI command clears the entire client blacklist:

```
(host) [md] #stm purge-blacklist-clients
```



These commands only remove the blacklisted clients from a particular managed device and not from the Mobility Master or other managed devices.

Working with WIP Advanced Features

Device Classification is the first step in securing the corporate environment from unauthorized wireless access. Adequate measures that quickly shut down intrusions are critical in protecting sensitive information and network resources. APs and stations must be accurately classified to determine whether they are valid, rogue, or a neighboring AP. Then, an automated response can be implemented to prevent possible intrusion attempts.

TotalWatch™ allows for detecting devices that are running on typical operational channels. Tarpit Shielding provides a better way of containing devices that are deemed unauthorized. Both of these features are discussed in the sections that follow.

- Configuring TotalWatch on page 621
- Administering TotalWatch on page 624
- Tarpit Shielding Overview on page 625
- Configuring Tarpit Shielding on page 625

Configuring TotalWatch

Aruba 802.11n APs and non-11n APs in AM-mode support for TotalWatch is the ability to scan all channels of the RF spectrum, including 2.4-and 5 GHz bands as well as the 4.9 GHz public safety band. TotalWatch also provides 5 MHz granular channel scanning of bands for rogue devices and dynamic scanning dwell times to focus on those channels with traffic. TotalWatch provides an advanced set of

features to detect unauthorized wireless devices and a set of customized rules are used to highlight devices that truly pose a threat to the network.



TotalWatch is supported on APs deployed in the AM-mode only.

TotalWatch provides monitoring support for the entire WLAN spectrum. Aruba APs in the AM-mode can monitor the following frequencies:

- 2412 MHz to 2472 MHz in the 2.5 GHz band.
- 5100 MHz to 5895 MHz in the 5 GHz band.

Aruba APs in AM-mode can scan the following additional frequencies:

- 2484 MHz and 4900 MHz to 5000 MHz (J-channels)
- 5000 to 5100 MHz

If the AP is HT-capable, these frequencies are scanned in the 40 MHz mode.

This section includes the following topics:

- Understanding TotalWatch Channel Types and Qualifiers on page 622
- Understanding TotalWatch Monitoring Features on page 622
- Understanding TotalWatch Scanning Spectrum Features on page 623
- Understanding TotalWatch Channel Dwell Time on page 623
- Understanding TotalWatch Channel Visiting on page 623
- Understanding TotalWatch Age out of Devices on page 624

Understanding TotalWatch Channel Types and Qualifiers

Based on the regulatory characteristics, channels are categorized into the following types:

- **Reg-domain Channels** : A channel that belongs to the regulatory domain of the country in which the AP is deployed. The set of channels that belong to this group is a subset of the channels in the all-reg-domain channel group.
- **All-reg-domain Channels** : A valid non-overlapping channel that is in the regulatory domain of at least one country. The channels in this category belong in the frequency ranges of:
  - 2412 MHz to 2472 MHz in the g-band.
  - 5100 MHz to 5895 MHz in the a-band.
- **Rare Channel** : Channels that fall into a frequency range outside of the regulatory domain; 2484 MHz and 4900 MHz-4995 MHz (J-channels), and 5000-5100 MHz. The channels in this group do not belong to any other group.

Each of these channel types can have an associated qualifier:

- **Active Channel** : This qualifier indicates that wireless activity is detected on this channel by the presence of an AP or other 802.11 activity such as a probe request.
- **DOS Channel** : A channel where wireless containment is active. This channel should belong to the country-code channel (regulatory domain).

Understanding TotalWatch Monitoring Features

TotalWatch enables monitoring of all channels including regulatory domain and rare channels. You can select one of the following scanning modes for each radio AP:

- scan only the channels that belong to the AP's regulatory domain
- scan channels that belong to all regulatory domains
- scan all channels

Understanding TotalWatch Scanning Spectrum Features

TotalWatch scans the following frequencies.

- G-band—2412 MHz to 2472 MHz
- J-band—2484 MHz and 4900-4995 MHz
- A-band—5000-5100 MHz to 5895 MHz

Table 113 list the frequency-to-channel mapping used by TotalWatch.

Table 113: *Frequency to Channel Mapping*

Frequency	Channel
2412 – 2472 MHz (in increments of 5 MHz)	1 - 13
2484 MHz	14
5100 – 5895 MHz (in increments of 5 MHz)	20 - 179
4900 – 4995 MHz (in increments of 5 MHz)	180 - 199
5000 – 5100 MHz	200 - 219

Understanding TotalWatch Channel Dwell Time

When an AP (in am-mode) visits a channel, the amount of time the AP stays on that channel is known as the dwell time. The channel dwell time is a variable value based on the following channel types.

- **dwell-time-active-channel**: For channels where there is wireless activity. Default setting is 500 ms.
- **dwell-time-reg-domain channel**: For channels that belong to the AP's regulatory domain group (reg-domain) with no wireless activity. The default setting is 250 ms.
- **dwell-time-other-reg-domain-channel**: For channels that belong to the all regulatory domain group (all-reg-domain) with no wireless activity The default setting is 250 ms.
- **dwell-time-rare-channel**: For channels in the rare group where no wireless activity is detected. The default value is 100 ms.

Use the **rf am-scan-profile** command to set the dwell time and scan mode.

Understanding TotalWatch Channel Visiting

The Active and DOS channels are visited more frequently than the other channels. The order of preference in selecting the next channel is:

1. DOS
2. Active
3. reg-domain
4. All-reg-domain
5. Rare

Once a channel is selected, the dwell time for that channel is determined based on the channel type. At the end of the dwell time, a new channel is picked.

Understanding TotalWatch Age out of Devices

ArubaOS uses a combination of inactivity time and unseen time to age out a device. This ensures that the channel is scanned a sufficient number of times before a device ages out. AM module maintains the following parameters:

- **Discovered Time**: The absolute time, in seconds, since the device was discovered.
- **Monitored Time**: The number of times the channel was scanned since discovery.
- **Inactivity Time**: The number of times the device was not seen when the channel is scanned.
- **Unseen Time**: The absolute time, in seconds, since the device was last seen.

Administering TotalWatch

The AM module will initialize the channel list for each of the AP's radio based on the scan mode setting for the radio. For example, if scan mode is set to rare, then the channel list will contain all possible channels. You can view these channels by using the **show ap arm scan-times** command.

The following sections provide information on configuring per radio settings, configuring per AP settings, and licensing:

Configuring Per Radio Settings

For each radio, you can configure the following settings (for detailed information on commands, refer to the*ArubaOS 8.7.0.0 Command Line Reference Guide*):

- the dwell times for the various channel types
- the channel list that should be used for scanning

These settings are configured via the command **rfam-scan-profile**, which can be attached to the two profiles, **dot11a-radio-profile** and **dot11g-radio-profile**.

The **am-scan-profile** includes the following parameters that can be configured:

- rf am-scan-profile <name>
- scan-mode [reg-domain | all-reg-domain | rare]

The default setting is the all-reg-domain. This is consistent with the default functioning of the AM scanning where the radio scans channels belonging to all regulatory domains.

Configuring Per AP Settings

If the AP is a dual-band single radio AP, an option is available to specify which band should be used for scanning in AM-mode. This setting is available in the **ap system-profile**, via the am-scan-rf-band command.

```
ap system-profile <name>
am-scan-rf-band [a | g | all]
```

The default value is "all", which is consistent with the prior behavior. This setting is ignored in the case of a dual radio AP.

There are four parameters that controls the age out of devices in the AM module.

```
ids general-profile <name>
ap-inactivity-timeout
sta-inactivity-timeout
ap-max-unseen-timeout
sta-max-unseen-timeout
```

The inactivity timeout is the number of times the device was not “seen” when the channel was scanned. The unseen timeout is the time, in seconds, since the device was last “seen.”

The **show ap monitor scan-info/channel** commands provide details of the channel types, dwell times, and the channel visit sequence.

```
(host) # show ap monitor scan-info ap-name rb-121
```

Licensing

The ability to perform rare scanning is available only with the RFprotect license. However, the AP can scan **reg-domain** or **all-reg-domain** channels without the RFprotect license.

Tarpit Shielding Overview

The Tarpit Shielding feature is a type of wireless containment. Detected devices that are classified as rogues are contained by forcing client association to a fake channel or BSSID. This method of tarpitting is more efficient than rogue containment via repeated de-authorization requests. Tarpit Sheilding works by spoofing frames from an AP to confuse a client about its association. The confused client assumes it is associated to the AP on a different (fake) channel than the channel that the AP is actually operating on, and will attempt to communicate with the AP in the fake channel.

Tarpit Shielding works in conjunction with the deauth wireless containment mechanism. The deauth mechanism triggers the client to generate probe request and subsequent association request frames. The AP then responds with probe response and association response frames. Once the monitoring AP sees these frames, it will spoof the probe-response and association response frames, and manipulates the content of the frames to confuse the client.

A station is determined to be in the Tarpit when we see it sending data frames in the fake channel. With some clients, the station remains in tarpit state until the user manually disables and re-enables the wireless interface.

Configuring Tarpit Shielding

Tarpit shielding is configured on an AP using one of two methods:

- **Disable all clients** : In this method, any client that attempts to associate with an AP marked for containment is sent spoofed frames.
- **Disable non-valid clients** : In this method, only non-authorized clients that attempt to associate with an AP are sent to the tarpit.

The choices for disabling Tarpit Shielding on an AP are:

- Deauth-wireless-containment
- Deauth-wireless-containment with tarpit-shielding (excluding-valid-clients)
- Deauth-wireless-containment with tarpit-shielding

This sections provide information on enabling Tarpit Shielding and licensing CLI commands.

Enabling Tarpit Shielding

The following CLI command configures **Tarpit Shielding** (for detailed information on commands refer to the *ArubaOS Command Line Reference Guide*).

```
(host) [mynode] (config) #ids general-profile default

(host) [mynode] (IDS General Profile "default") #wireless-containment [deauth-only | none
| tarpit-all-sta | tarpit-non-valid-sta]
```

The following CLI command displays the updated Tarpit Shielding status and the spoofed frames generated for an AP:

- **show ap monitor stats**
- **show ap monitor containment-info**

Understanding Tarpit Shielding Licensing CLI Commands

Under the **ids general-profile default wireless-containment** command, the **tarpit-non-valid-sta** and **tarpit-all-sta** options are available only with a RFprotect license. The **deauth-only** and **none** options are available with the Base OS license.

This chapter gives an overview of the basic functions of APs, and describes the process to install and configure the APs on your network. When an AP is first installed on the network and powered on, the AP locates the managed device and the AP's designated configuration is sent from the managed device to the AP.



APs cannot terminate either on Mobility Master or a master controller. They must terminate on managed devices only.

The default management credentials for IAP and UAP for WebUI, SSH, and console access are:

- **Username**: admin
- **Password**: serial number of the AP

The same credentials will be used if IAPs running ArubaOS versions prior to ArubaOS 8.5.0.0 are upgraded to ArubaOS 8.5.0.0 and factory reset. If the IAP is part of a cluster, the username will be admin and the password will be the serial number of any of the APs in the cluster.

If the IAP is running software version prior to ArubaOS 8.5.0.0, **admin** would continue to be the default password.

Before Deploying an AP

Before you install APs in a network environment, you must ensure that the APs are able to locate and connect to the managed device. Specifically, you must configure firewall settings to allow APs to obtain software images and configuration settings from the managed device. You must also verify that the APs are able to locate the Mobility Master, and verify that each AP is assigned a valid IP address when connected to the network.



Mobility Master cannot be used as an AP master since APs are not allowed to terminate on a Mobility Master. If the AP manager on Mobility Master receives an AP HELLO message, the message is dropped.

The following topics describe the pre-deployment tasks. Click any of the following links for more information:

- Controller Licenses on page 628
- Understanding Firewall Port Configuration in Aruba Devices on page 831
- Discovery of Controller on page 629
- Duplicate Address Detection on page 633
- AP Provisioning on page 633

Mesh AP Pre-configuration

Mesh APs require the following additional steps to define the mesh networking environment:

- Configuring Mesh Cluster Profiles on page 763
- Creating and Editing Mesh Radio Profiles on page 752

Remote AP Pre-configuration

Remote APs require the following additional step to identify valid APs in the remote AP whitelist:

- Configuring Certificate Remote AP on page 848

Controller Licenses

ArubaOS supports a centralized licensing architecture, which allows a group of managed devices to share a pool of licenses. A primary and backup Mobility Master can share a single set of licenses, eliminating the need for a redundant license set on the backup server. For information on license types, usage, and license installation, see *ArubaOS Licensing Guide*.

Understanding Firewall Port Configuration in Aruba Devices

This section describes the network ports that need to be configured on the firewall to allow proper operation of the network.

Communication Between Managed Devices

Configure the following ports to enable communication between any two managed devices:

- IPsec (UDP port 4500) for communication between Mobility Master and a managed device.
- IPsec (UDP ports 500 and 4500) and ESP (protocol 50). PAPI between Mobility Master and a managed device 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 managed device
- IKE (UDP 500)
- ESP (protocol 50)
- NAT-T (UDP 4500)

Communication Between APs and the Managed Device

APs use Trivial File Transfer Protocol (TFTP) during their initial boot to grab their software image and configuration from the managed device. After the initial boot, the APs use FTP to retrieve their software images and configurations from the managed device. In many deployment scenarios, an external firewall is situated between various Aruba devices.

Configure the following ports to enable communication between an AP and the managed device:

- PAPI (UDP port 8211). If the AP uses DNS to discover the LMS managed device, the AP first attempts to connect to the managed device. (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 managed device.
- FTP (TCP port 21)
- TFTP (UDP port 69). All campus APs; If there is no local image on the AP or if the image needs to be upgraded (for example, a new AP), the AP will use TFTP to retrieve the initial image. For remote APs, upgrade the image only by FTP and not TFTP.
- SYSLOG (UDP port 514)

- PAPI (UDP port 8211)
- GRE (protocol 47)
- Control Plane Security (CPsec) uses UDP port 4500

Communication Between Remote APs and the Managed Device

Configure the following ports to enable communication between a Remote AP (IPsec) and a managed device:

- NAT-T (UDP port 4500)
- TFTP (UDP port 69)



TFTP is not needed for normal operation. If the remote AP loses its local image for any reason, it will use TFTP to download the latest image.



In an IPv6 deployment, you must configure UDP port 500 on the firewall to establish an IPv6 IPsec connection for Remote APs.

Discovery of Controller

An AP can discover the IP address of the controller from a DNS server, from a DHCP server, or using the Aruba Discovery Protocol.

At boot time, the AP builds a list of managed device IP addresses and then tries these addresses in order until it successfully reaches a managed device. The AP constructs its list of managed device addresses as follows:

- If theprovisioning parameter is set to a DNS name, that name is resolved and all resulting addresses are put on the list. If it is set to an IP address, that address is put on the list.
- If the provisioning parameter is not set and a managed device address was received in DHCP Option 43, that address is put on the list.
- If the provisioning parameter is not set and no address was received via DHCP option 43, ADP is used to discover a managed device address and that address is put on the list.
- Managed device addresses derived from the **server-name** and **server-ip** provisioning parameters and the default managed device 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 IP addresses provides an enhanced redundancy scheme for managed device that are located in multiple data centers separated across Layer-3 networks.

Controller Discovery Using DNS

When using DNS, AP learns multiple IP addresses to associate with a managed device. If the primary node is unavailable or does not respond, the AP continues through the list of learned IP addresses until it establishes a connection with an available managed device. This takes approximately 3.5 minutes per managed device.



It is recommended you use a DNS server to provide APs with the IP address of the managed device because it involves minimal changes to the network and provides the greatest flexibility in the placement of APs.

APs are factory-configured to use the host name **aruba-master** for the managed device that terminates the APs. For the DNS server to resolve this host name to the IP address of the managed device, configure an entry on the DNS server for the name **aruba-master**.

Controller Discovery Using Aruba Discovery Protocol

ADP is enabled by default on all Aruba APs and managed devices. With ADP, APs send out periodic multicast and broadcast queries to locate the Mobility Master. ADP requires that all APs and managed devices are connected to the same Layer-2 network. If the devices are on different networks, you must use a Layer-3 compatible discovery mechanism, such as DNS, DHCP, or IGMP forwarding.

To use ADP discovery:

1. Execute the command **show adp config** to verify that ADP and IGMP join options are enabled on the managed device, If ADP is not enabled, you can re-enable ADP using the command **adp discovery enable** and **adp igmp-join enable**.
2. If the APs are not in the same broadcast domain as the Mobility Master, you enable multicast on the network (ADP multicast queries are sent to the IP multicast group address 239.0.82.11) for the Mobility Master to respond to the APs' queries. Ensure that all routers are configured to listen for IGMP join requests from the controller and can route these multicast packets.

Controller Discovery Using a DHCP Server

You can configure a DHCP server to provide the IP address or VRRP IP address of the Mobility Controller. Configure the DHCP server to send the managed device's IP address using the DHCP vendor-specific attribute option 43. The APs identify themselves with a vendor class identifier set to **ArubaAP** in their DHCP requests. When the DHCP server responds to a request, it will send the managed device'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 managed device provided by DHCP is not available, the AP can use the other IP addresses provisioned or learned by DNS to establish a connection. For more information on how to configure vendor-specific information on a DHCP server, see DHCP with Vendor-Specific Options on page 1255 or refer to the documentation included with your server.

Enhancements to AP Master Discovery

Starting from ArubaOS 8.7.0.0, users can configure the preferred IP protocol for AP master discovery. The default IP protocol for AP master discovery is IPv4. Follow one of the procedures below to configure the preferred IP protocol:

- While deploying a new AP, select **IPv4** or **IPv6** for **Controller discovery preference** in the **Task > AP Settings** page.
- Navigate to **Configuration > Access Points**. Select an AP for which you need to configure the preferred IP protocol. Under **General**, select **IPv4** or **IPv6** for **Controller discovery preference**.
- Navigate to **Configuration > System > Profiles > AP > Provisioning**. Select a profile and select**IPv4** or **IPv6** from the **Master Preference** drop-down list.

The following CLI commands configure the preferred IP protocol for an AP provisioning profile:

```
(host) [mynode] (config) #ap provisioning-profile test
(host) [mynode] (Provisioning profile "test") #master-preference ipv6
```

The following CLI commands configure the preferred IP protocol using the **provision-ap** command:

```
(host) [mynode] (config) #provision-ap
```

```
(host) [mynode] (config-submode)#master-preference ipv6
```



If a static IP address is already configured as the master IP address, the preferred IP protocol will not take effect.

Enable DHCP to Provide APs with IP Addresses

Each AP requires a unique IP address on a subnetwork that has connectivity to a managed device. It is recommended you use the DHCP to provide IP addresses for APs; the DHCP server can be an existing network server or a managed device configured as a DHCP server. The APs can avoid IP conflicts due to IP misconfiguration or DHCP server malfunctioning. You can now identify and debug IP conflict issue in the local network.



If you do not enable DHCP, each AP must be manually configured with an IP address through the AP provisioning profile.

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.)

The Mobility Master can configure the managed device as a DHCP server to assign an IP address to the AP. The managed device must be the only DHCP server for this subnetwork.



When APs with more than one Ethernet interface are not under a managed device, the APs act as DHCP servers to wired clients. This occurs when both the Ethernet ports of the APs are connected to the uplink switch, without configuring LACP on the uplink switch.

The following procedure describes how to configure the DHCP services:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > DHCP** tab.
2. Expand **DHCP Server.**
3. Select the check box for either **IPv4 DHCP server** or **IPv6 DHCP server**.
4. In the **Pool Configuration** table, click **+**.
5. Enter information about the subnetwork for which IP addresses are to be assigned. The parameters are described in Table 114.
6. Click **Submit**.

Table 114: *DHCP Configuration Parameters*

Parameter	Description
IP Version	Select the IP version used by the DHCP pool configuration
Pool name	Enter the name of the DHCP pool.
Default routers	Enter an IP address to assign the DHCP default router for the managed device
DNS Servers	Enter an IP address to assign the DHCP DNS server for the managed device

Parameter	Description
Import from DHCP/PPPoE	Enables/disables importing of DNS server configurations.
Domain name	Domain name used by the client.
WINS	Enter IP address to assign Windows Internet Name Service servers. When entering multiple servers, each server must be separated by a space.
Import from DHCP/PPPoE	Enables/disables importing WINS server configurations.
Lease days	Length of time a device may lease the DHCP in days. Entering **0** indicates no time limit.
Lease hours	Length of time a device may lease the DHCP in hours.
Lease minutes	Length of time a device may lease the DHCP in minutes.
Lease seconds	Length of time a device may lease the DHCP in seconds.
Network IP address	Subnetwork for the pool
Network IP mask	Enter an IP address to assign the netmask for the DHCP pool
Option	DHCP option number

Excluding IP Addresses

If there are addresses that should not be assigned in the subnetwork:

1. Click **+** in the **Excluded Address Range** table that corresponds with the IP version used.
2. Enter the address range in the **Add Excluded Address** section.
3. Click **Submit**.

The following CLI commands configure the DHCP services:

```
(host)[node](config)# ap system-profile <profile>
rap-dhcp-default-router <ip-addr>
rap-dhcp-dns-server <ip-addr>
rap-dhcp-lease
rap-dhcp-pool-end
rap-dhcp-pool-start
rap-dhcp-pool-netmask
rap-dhcp-server-id
rap-dhcp-server-vlan
(host)[node](config)# ip dhcp
adaptive
default-pool
excluded-address
load-balance
ping-check
pool
(host)[node](config)# service
dhcp
dhcpv6
```

IP Conflict Detection

An IP conflict occurs when the same IP address is assigned to another host on the network. This causes an AP to lose connectivity to the local network. Starting from ArubaOS 8.4.0.0, an AP can detect and resolve an IP conflict in the network using a Gratuitous ARP probe.

Starting from ArubaOS 8.7.0.0, Duplicate Address Detection (DAD) feature is supported for APs in IPv6 deployments.

DAD is an IP conflict detection mechanism that identifies IP conflicts when the AP obtains IPv6 address from static, RA, or DHCPv6 server. However, DAD prevents IP conflicts only when the AP obtains IPv6 address from DHCPv6 server. IPv6 conflict usually occurs when two hosts on the same network use the same IPv6 address, which creates unstable connection on both the hosts.

When a new device joins and requests for an IPv6 address, the DHCPv6 server allocates a free IPv6 address from the IP pool. In this process, if the DHCPv6 server malfunctions and assigns multiple devices with a single IPv6 address, or if the AP obtains an IPv6 address from DHCPv6 server that is already allocated to a device in the same network, an IP conflict occurs. For example, if the AP receives 2001:1234::abcd as the IPv6 address from DHCPv6 server during boot process, and another host on the VLAN also uses 2001:1234::abcd IPv6 address, then DAD identifies the duplicate IPv6 address conflict and renews a free IPv6 address from the DHCPv6 server. DAD verifies if a configured unicast IPv6 address is unique and has not been used anywhere else before the unicast IPv6 address is assigned to a VLAN interface on the AP. If the unicast IPv6 address is a duplicate, the address is not used in the deployment.



DAD is enabled and configured on the default IPv6 configuration during AP boot process.

AP Provisioning

AP provisioning settings allow you to define a set of additional provisioning information for an AP, such as USB modem settings, PPPoE values, or configuration settings to provision an AP as a Remote AP.

Ensure that any provisioning changes you make are complete and accurate before you save those settings. If an AP is configured incorrectly with erroneous parameters, that AP may be lost. If you want to provision APs with more than one interface, you can also configure the USB settings and interface priority levels using an AP provisioning profile.

The following procedure describes how to provision APs.

1. Navigate to the **Configuration > Access Points** window.
2. Select the AP to which you want to add new provisioning settings, then click **Provision**. The AP provisioning settings are divided into two groups. By default, the ArubaOS WebUI displays configuration settings described in Table 115.

Table 115: *AP Provisioning Profile Parameters*

Parameter	Description
Name	Name assigned to an AP. An AP requires a reboot before a new AP name takes effect. Therefore, wait until there is little or no client traffic passing through the AP before renaming it.
AP Group	AP group to which the AP is assigned.
Remote-AP	Select this check box to provision the APs as Remote APs. If you are provisioning Remote APs, you must also add the remote APs to the Remote AP whitelist. For details, see Remote Access Points on page 844.

Parameter	Description
Controller discovery	Select **Use AP discovery protocol (ADP)** if you want to provide the AP with its managed device IP address, or select **Static** to manually define the managed device IP for that AP. If you select the **Static** option, you are prompted to enter the managed device's DNS name or IP address. ADP is enabled by default on all Aruba APs and managed devices. With ADP, APs send out periodic multicast and broadcast queries to locate the Mobility Master. ADP requires that all APs and managed devices are connected to the same Layer-2 network. If the devices are on different networks, you must use a Layer-3 compatible discovery mechanism, such as DNS, DHCP, or IGMP forwarding.
Controller discovery preference	Select **IPv4** or **IPv6** depending on your preference. The default is **IPv4**.
IP	Select **DHCP** if you have configured a DHCP server to provide the AP with the AP IP address, or select **Static** to manually define the AP IP address. If you select the **Static** option, you are prompted to enter the following information for the selected AP: ▪ IPv4 address, netmask, internet gateway used by the AP, and DNS server. ▪ IPv6 address, netmask, internet gateway used by the AP, and DNS server.
TFTP Server (Select **Show advanced options**)	IPv4 / IPv6 address of the TFTP server from which the AP can download its boot image.
Coverage Area	This setting defines the type of installation (**indoor** or **outdoor**). The **default** option indicates that the installation mode is determined by the AP model type.
Single Chain Mode	If this option is enabled for an 802.11n-capable radio, the radio will operate in single-chain mode, and will transmit and receive data using only legacy rates and single-stream HT rates up to MCS 7. This parameter is disabled by default.
PEAP username	Username of AP so that AP can authenticate to 802.1X using PEAP.
PEAP password	Password of AP so that AP can authenticate to 802.1X using PEAP.
EAP-TLS	Enable AP to 802.1x using EAP-TLS.
EAP-TLS use factory certificate	Enable AP to use factory certificates when doing 802.1x EAP-TLS.



It is recommended not to connect both the Ethernet ports of the APs to the uplink switch, because the APs act as DHCP servers to wired clients when LACP is not configured on the uplink switch. This occurs when APs with more than one Ethernet interface are not under a managed device.

Auto-Provisioning of APs

AP Auto-Provisioning settings allow you to automate and simplify AP provisioning by assigning pre-provisioning rules to new APs. When new APs are connected to the network, the pre-provisioning rules are applied and the APs are automatically provisioned based on the conditions and actions defined in the rules. This enables the managed device to do bulk provisioning of APs with different attributes.

Each rule is assigned a priority level to determine the precedence of the rules. When an unprovisioned AP connects to a network, the AP checks the rules that have the highest priority. If the AP matches the conditions of the rule, the AP sets each action to the corresponding provision parameter and provisions itself. If the AP does not match the pre-defined criteria of the rule, the AP checks for the next priority rule that matches the condition.



You can apply auto-provisioning of APs only on Mobility Master and stand-alone controllers. You cannot apply provisioning rules to APs that have already been configured by whitelist, manual provisioning, or provisioning profile methods.

Configuring Provisioning Rules to APs

The following procedure describes how to configure the provisioning rules to an AP:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Access Points > Provisioning Rules** tab.
2. Under **AP Provisioning Rules**, click **+** to create a new provisioning rule.
3. To edit an existing rule, click the provisioning rule entry from the **Provisioning Rules** table. To delete a existing rule, click the trash icon on the right.). The **New Provisioning Rule** window appears.
4. Configure the **New Provisioning Rule** settings, then click **Submit**.

Table 116 describes the new provisioning rule parameters.

Table 116: *New Provisioning Rule Parameters*

Parameter	Description
Name	Enter a provisioning rule name.
Conditions	Specify conditions to narrow the scale of AP based on each criteria in the conditions.
Attribute	Select one of the following options from the drop-down list: ■ **IP Range**: Specify the IPv4 or IPv6 address range to check if the IP address of the AP is within this range. ■ **Network Address**: Enter an IPv4 or IPv6 network address along with netmask to check if the IP address of the AP is within this network address. The network address includes the prefix (length of the netmask) in the following format: x.x.x.x/prefix or x:x::x/prefix. **NOTE:** The allowed values of the netmask prefix are as follows: ■ 0—32 for IPv4 network address; 0—128 for IPv6 network address. • **AP Model**: Select an AP model from the drop-down list to apply the rule to this particular AP. • **Any AP**: Select this option if you want to apply the provisioning rule's condition to any AP. ■ You must drag the rules up or down to reorder or reassign the priority of the rules in the WebUI.

Parameter	Description
	▪ The rule with **Any AP** condition has the lowest priority and is applied only when the previous conditions are not met. You cannot move the position of the **Any AP** rule in the WebUI. The **Any AP** condition is unique and is exclusive of other conditions.
Actions	Specify actions that corresponds to the conditions. When an AP satisfies all the criteria in conditions, the AP executes the actions applied by the rule.
Attribute	Select one of the following options from the drop-down list: ▪ **Assign to AP Group**: From the drop-down list, select the AP group that you want to assign to the AP. ▪ **Set Antenna Gain for Dual 5GHz mode**: Set values for Radio 0 and Radio 1 for APs that support Dual 5 GHz mode. ▪ **Set Antenna Gain for Dual Band mode**: Set values for 2.4GHz and 5GHz antenna for APs that support both the bands. **NOTE:** ▪ The Antenna gain attributes are applicable only to the APs with external antenna. ▪ You cannot configure the AP group that is set to **default**.

The following CLI commands add the condition and associated parameter for auto-provisioning rule:

```
(host) [mm] (config) #ap provisioning-rule test
(host) [mm] (ap provisioning rule "test") #condition ap-type AP-103
```

To add the action and associated parameter for auto-provisioning rule:

```
(host) [mm] (ap provisioning rule "test") #action ap-group corp1
```

To display the provisioning rule for an AP:

```
(host) [mm] (config) #show ap provisioning-rule test
ap provisioning rule "test"
---------------------------
Parameter                                                  Value
---------                                                  -----
ip range                                                   N/A
network                                                    3.3.3.3/1
any AP                                                     false
AP Type                                                    AP-UNKNOWN
ap group                                                   N/A
Antenna gain for 802.11g                                   N/A
Antenna gain for 802.11a                                   N/A
Radio 0 5GHz Antenna gain for APs support Dual 5GHz mode   N/A
Radio 1 5GHz Antenna gain for APs support Dual 5GHz mode   N/A
```

To select the required rules from provisioning rule for auto-provisioning and define their priority:

```
(host) [mm] (config) #ap provisioning-rules
(host) [mm] (ap provisioning rules) #provision-rule test priority 20
(host) ^[mm] (ap provisioning rules) #write memory
```

To display the rules based on the priority level:

```
(host) [mm] (ap provisioning rules) #show provisioning-rule-info summary
Auto provision Rule Info
------------------------
Rule Name Priority Hit times Success count
--------- -------- --------- -------------
ap324       1        0          0
ip36        3        0          0
network     5        0          0
ip46        7        0          0
```

To delete the provisioning rule:

```
(host) [mm] (config) #no ap provisioning-rule test
```

- The Mobility Master or standalone controller checks the list of unprovisioned APs on an LMS every 10 seconds for a total of six times, and provisions the AP that matches the rule. If the provisioning still fails after six attempts, the AP is ignored and is set to unprovisioned state. If the AP is provisioned by other provisioning methods within the 10 seconds interval, the AP is removed from the list and is not provisioned automatically.
- When you manually enable the CPsec and add the whitelist of the AP with ap-name or ap-group configured, the AP is not provisioned automatically.
- You can create a maximum of 32 rules for an AP group.
- There is no action in the rule to provision AP to remote AP or mesh AP. Hence, you must provision it manually or set the remote AP by using provisioning profile.
- You can manually provision APs before or after their auto-provisioning, to set the extra provisioning parameter to the APs.

Overview of Basic Functions of an AP

APs use AI-powered RF optimization, rich user and application intelligence, and smart management options to improve user experiences, enhance Quality of Service (QoS), and support digital workplace initiatives. This section describes the basic functionalities of an AP. Use the Mobility Master WebUI and command-line interface to configure APs.

Table 117: *AP Configuration Function Overview*

Features and Function	Description
WLANs	A WLAN permits wireless clients to connect to the network. An AP broadcasts the SSID (which corresponds to a WLAN configured on the Mobility Master to wireless clients. APs support multiple SSIDs. WLAN configuration includes the authentication method and the authentication servers by which wireless users are validated for access. The WebUI includes a WLAN Wizard that provides easy-to-follow steps to configure a new WLAN. **NOTE:** All new WLANs are associated with the ap-group named “default”.
AP operation	An AP can function as an AP that serves clients, as an AM performing network and RF monitoring, or as a hybrid AP that serves both clients and performs spectrum analysis a single radio channel. You can also specify the regulatory domain (the country) which determines the 802.11 transmission spectrum in which the AP will operate. Within the regulated transmission spectrum, you can configure 802.11a, 802.11b/g, or 802.11n (high-throughput) radio settings. **NOTE:** The 802.11n features, such as high-throughput and 40 MHz configuration settings, are supported on APs that are 802.11n standard compliant.
Quality of Service (QoS)	Configure Voice over IP call admission control options and bandwidth allocation for 5 GHz (802.11a) or 2.4 GHz (802.11b/g) frequency bands of traffic.
RF Management	Configure settings for balancing wireless traffic across APs, detect holes in radio coverage, or other metrics that can indicate interference and potential problems on the wireless network. 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.

Features and Function	Description
Intrusion Detection System	Configure settings to detect and disable rogue APs, adhoc networks, and unauthorized devices, and prevent attacks on the network. You can also configure signatures to detect and prevent intrusions and attacks.
Mesh	Configure Aruba APs as mesh nodes to bridge multiple Ethernet LANs or extend wireless coverage. A mesh node is either ■ mesh portal: an AP that uses its wired interface to reach the managed device ■ mesh point: an AP that establishes a path to the managed device via the mesh portal ■ mesh auto: an AP that automatically detects the mesh role and configures mesh portal or mesh point. **NOTE:** ■ Starting from ArubaOS 8.4.0.0, you can set mesh role to **auto** under AP provisioning. Mesh auto enables auto-detection of mesh role based on system initialization or operation. The role switches between mesh point or mesh portal depending on the ethernet link and mesh role detected packets. ■ Mesh environments use a wireless backhaul to carry traffic between mesh nodes. This allows one 802.11 radio to carry traditional WLAN services to clients and one 802.11 radio to carry mesh traffic and WLAN services. Secure Enterprise Mesh on page 738 contains more specific information on the Mesh feature.
AP Boot Time	An AP can take up to a minute to boot up currently. The fast boot feature enables the APs to boot up within less time. Faster boot time results in an AP boot time from boot to shell prompt within 1 minute. **NOTE:** AP fast boot is supported on AP-534, AP-535, and AP-555 access points only.

AP Configuration Profiles

An AP configuration profile is a general name to describe any of the different groups of settings that can be defined, saved, and applied to an Access Point. ArubaOS has many different types of profiles that each allow you to configure a different aspect of an overall configuration of an AP. ArubaOS also contains a predefined "default" profile for each profile type. You can use the predefined settings in these default profiles, or create entirely new profiles that you can edit as required.

Each different AP configuration profile type can be managed using the CLI or the WebUI. To see a full list of available configuration profiles using the command-line interface, access the CLI and issue the command **show profile-hierarchy**.

To view available configuration profiles using the WebUI, navigate to **Configuration > System**, then select the **Profiles** tab.



The profile types that appear in the **All Profiles** list may vary, depending upon the controller configuration and available licenses.

The following sections provide information on AP profiles and RF Management profiles:

- AP Profiles on page 640
- RF Management Profiles on page 641

AP Profiles

The following AP profiles configure AP operation parameters, radio settings, port operations, regulatory domain, and SNMP information:

- **AM Filter:** Clients may assign APs or AP groups to AM filter profiles. These profiles collect data that is used to identify and monitor APs, wireless clients, and mesh nodes within the network.
- **AP Authorization:** Allows you to assign a provisioned but unauthorized AP to a AP group with a restricted configuration profile. For details see Configuring Remote AP Authorization Profiles on page 872.
- **AP Ethernet Link**: Sets the duplex mode and speed of the AP's Ethernet link. The configurable speed is dependent on the port type, and you can define a separate Ethernet Interface profile for each Ethernet link. For details on configuring this profile, see Table 130.
- **AP LACP LMS map information:** Maps a LMS IP address to a GRE striping IP address. If the AP fails over to a standby or backup Mobility Master, the AP LACP LMS map information profile on the new LC defines the striping IP address that the AP uses for link aggregation. For details, see Configuring Port Channel LACP on page 119.
- **AP LLDP and AP LLDP-MED Network Policy:** LLDP is a Layer 2 protocol that allows network devices to advertise their identity and capabilities on a LAN. The LLDP-MED Network Policy profile defines the VLAN, priority levels, and DSCP values used by a voice or video application. Wired interfaces on Aruba APs support LLDP by periodically transmitting LLDP Protocol Data Units comprised of selected TLV elements. The AP LLDP profile identifies which TLVs will be sent by the AP. For details, see Understanding Extended Voice and Video Features on page 1124.
- **AP MultiZone:** The MultiZone feature allows an AP to terminate to multiple managed devices that reside in different zones. A zone is a collection of managed devices under a single administration domain. For details, see MultiZone on page 453.
- **AP system:** Defines administrative options for the managed device, including the IP addresses of the local, backup, and master controllers, RTLS server values and the number of consecutive missed heartbeats on a GRE tunnel before an AP reboots. For details on configuring this profile, see Optional AP Configuration Settings on page 708.
- **AP Wired Port:** Specifies a AAA profile for users connected to the wired port on an AP.

- **Dump Collection:**Specifies the profile for collecting core dump when an AP process crashes. For details, see Configuring the Dump Collection Profile.
- **EDCA Parameters (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 Working with QoS for Voice and Video on page 1083.
- **EDCA Parameters (Station):** Client-to-AP traffic prioritization parameters, including EDCA parameters for background, best-effort, voice and video queues. For additional information on configuring this profile, see Working with QoS for Voice and Video on page 1083.
- **Regulatory Domain:** Defines the AP's country code and valid channels for both legacy and high-throughput 802.11a and 802.11b/g radios.
- **Spectrum Local Override:** configure an individual AP radio as a spectrum monitor, For details, see Converting AP to Spectrum Monitor on page 906.
- **Wi-Fi Uplink:** Configure a Wi-Fi uplink profile that allows an AP running ArubaOS to connect to an external wireless network or a managed device by using a third-party AP, such as a Mi-Fi device. For details on configuring this profile, see Uplink Monitoring and Load Balancing on page 247.
- **Wired AP**: Determines if 802.11 frames are tunneled to the managed device using GRE tunnels, bridged into the local Ethernet LAN, or configured for a combination of the two (split-mode). In tunnel forwarding mode, the AP handles all 802.11 association requests and responses, but sends all 802.11 data packets, action frames and EAPOL frames over a GRE tunnel to the managed device for processing. When a remote AP or campus AP is in bridge mode, the AP handles all 802.11 association requests and responses, encryption/decryption processes, and firewall enforcement. In split-tunnel mode, 802.11 frames are either tunneled or bridged, depending on the destination (corporate traffic goes to the managed device, and Internet access remains local). For details, see Configuring Ethernet Ports for Mesh on page 770

RF Management Profiles

The following profiles configure radio tuning and calibration, AP load balancing, and RSSI metrics:

- **802.11a:** defines AP radio settings for the 5 Ghz frequency band, including the ARM profile and the high-throughput (802.11n) radio profile. For additional information on configuring this profile, see 2.4 GHz and 5 GHz Radio RF Management on page 676.
- **802.11g:** defines AP radio settings for the 2.4 Ghz frequency band, including the ARM profile and the high-throughput (802.11n) radio profile. Each 802.11a and 802.11b radio profile includes a reference to an ARM profile.
- If you want ARM to dynamically select the best channel and transmission power for the radio, verify that the 802.11a/802.11g radio profile references an active and enabled ARM profile. If you want to manually select a channel for each AP group, create separate 802.11a and 802.11g profiles for each AP group and assign a different transmission channel for each profile. For additional information on configuring this profile, see 2.4 GHz and 5 GHz Radio RF Management on page 676.
- **Adaptive Radio Management:** defines the 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 Configuring ARM Profiles on page 588.
- **AM Scanning**: Aruba 802.11n APs and non-11n APs in AM-mode support the TotalWatch scanning feature giving them the ability to scan all channels of the RF spectrum, including 2.4-and 5-Ghz bands as well as the 4.9-GHz public safety band. The AM Scanning profile enables this feature, and defines the dwell types for different channel types.

- **High-throughput radio:** manages high-throughput (802.11n) radio settings for 802.11n-capable APs. A high-throughput profile determines 40 Mhz tolerance settings, and controls whether or not the APs using this profile will advertise intolerance of 40 Mhz operation. (This option is disabled by default, allowing 40 MHz operation.) For additional information on configuring this profile, see High-Throughput APs on page 688.
- **RF Event Thresholds:** defines error event conditions, based on a customizable percentage of low-speed frames, non-unicast frames, or fragmented, retry or error frames. For additional information on configuring this profile, see 2.4 GHz and 5 GHz Radio RF Management on page 676.
- **RF Optimization:** 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 RSSI metrics.

Converting APs to Instant APs

Starting from ArubaOS 8.6.0.0, you can convert a Campus AP or a Remote AP to an Instant AP that is managed by Aruba Central, by using a new command—**ap convert**. However, Aruba does not support this feature for Instant AP deployments that are managed through AirWave or local WebUI, and recommends using this command only in lab or test environments for such deployments.

You can convert the APs, AP lists, or AP groups using local-flash or local image server options like ftp, tftp, http, https, or scp by copying the downloaded image from Aruba support to the local ftp/tftp/scp server. From that server, the managed device downloads the image to its ftp or tftp folder and then distributes the ftp or tftp URLs to Campus APs.



The converted APs are limited to the highest supported version of the corresponding Instant APs. For example, if IAP-225 runs only up to Aruba Instant 8.6.0.0 version, the converted AP-225 will also support up to Aruba Instant 8.6.0.0 version.

This feature also supports conversion of APs based on AP groups or AP lists, which allows the user to manage the conversion seamlessly and also, avoid the high load on a managed device.



Ensure to disable the load balancing feature in a cluster to avoid the AP's movement to different managed devices during conversion.

To convert APs using local-flash option, upload the images in flash before executing the following commands:

```
(host) [mynode] #ap convert active specific-aps local-flash <images>
(host) [mynode] #ap convert active all-aps local-flash <images>
```

To convert APs using image servers, execute one of the following commands depending on the mode:

```
(host) [mynode] #ap convert all-aps server ftp: <ftphost> <user> <images >
(host) [mynode] #ap convert specific-aps server ftp: <ftphost> <user> <images>
(host) [mynode] #ap convert all-aps server scp: <scphost> <user> <images >
(host) [mynode] #ap convert specific-aps server scp: <scphost> <user> <images>
(host) [mynode] #ap convert all-aps server tftp: <tftphost> <images >
(host) [mynode] #ap convert specific-aps server tftp: <tftphost> <images>
```

To add specific AP groups or AP names to convert, execute the following command:

```
(host) [mynode] #ap convert add ap-group <ap-group>
(host) [mynode] #ap convert add ap-name <ap-name>
```

To remove specific AP groups or AP names from list of conversion, execute the following command:

```
(host) [mynode] #ap convert delete ap-group <ap-group>
(host) [mynode] #ap convert delete ap-name <ap-name>
```

To clear all the APs from the list of conversion, execute the following command:

```
(host) [mynode] #ap convert clear-all
```

To abort the conversion of APs:

```
(host) [mynode] #ap convert cancel
```

Configuring Installed APs

APs and AMs are designed to require only minimal setup to make them operational in a user-centric network. Once APs have established communication with the managed device, apply advanced configuration to individual APs or groups of APs in the network using the WebUI on the managed device.

You can either connect the AP directly to a port on the managed device, or connect the AP to another switch or router that has layer-2 or layer-3 connectivity to the managed device. If the Ethernet port on the managed device is an 802.3af PoE port, the AP automatically uses it to power up. If a PoE port is not available, get an AC adapter for the AP. For more information, see the *Installation Guide* for the specific AP.



It is recommended not to connect both the Ethernet ports of the APs to the uplink switch, because the APs act as DHCP servers to wired clients when LACP is not configured on the uplink switch. This occurs when APs with more than one Ethernet interface are not under a managed device.

If you are configuring a new AP that has never been provisioned before, first connect the AP to the managed device according the instructions included with that AP. If you are re-provisioning or reconfiguring existing active APs, this step is not necessary, as the APs are already communicating with the managed device.

You can configure an AP using the AP wizard, the provisioning profile in the WebUI, or the managed device command-line interface. The individual configuration steps vary, depending on whether the AP is deployed as a Campus AP, Remote AP, or a Mesh AP.

This following sections describe the procedure to configure an installed AP with the basic settings it requires to become operational on the network:

- Configuring an AP using AP Wizard
- Configuring a Remote AP
- Verifying the Configuration

Configuring an AP using AP Wizard

The easiest way to provision any AP is to use the AP Wizard in the managed device WebUI. This wizard will walk you through the specific steps required to provision a Campus, Remote or Mesh AP. The Wizard includes a help tab that further describes each of the configuration tasks for that deployment type.

The following procedure describes how to access the AP wizard to provision an AP:

1. Select the managed device to which the AP will be provisioned.
2. Navigate to the **Configuration** > **Access Points** page.
3. Select the new AP from the **Campus APs** list, then click **Provision.**

4. In the **General** section, click the **AP Group** drop-down list and select the AP group to which this AP should be assigned. The AP group must have at least one virtual AP.
5. (Optional) Some AP models support an external antenna in addition to their internal antenna. If the AP you are provisioning supports an external antenna, the Provisioning window displays an additional **Antenna Parameters** section.
6. (Optional) To allow the remote AP to use PEAP to authenticate to 802.1X networks, select **Show Advanced Options** under the **General** tab, then enter a user name and password in the 802.1X Parameter using PEAP section.
7. In the **IP Settings** section, define how the AP should obtain its IP address. If you have configured a DHCP server to allow APs to get addresses using DHCP, select **Obtain IP address using DHCP**. For more information on configuring a DHCP server, see Duplicate Address Detection on page 633. Otherwise, select **Use the Following IP address** and enter the appropriate values in the following fields:
   - **IP address**: IP address for the AP, in dotted-decimal format
   - **Subnet mask**: Subnet mask for the IP, in dotted-decimal format.
   - **Gateway IP address**: The IP address the AP uses to reach other networks.
   - **DNS IP address**: The IP address of the Domain Name Server.
   - **Domain name**: (optional) The default domain name.
8. (Optional) Access points can be configured in single-chain mode, allowing the radios of those APs to transmit and receive data using only legacy rates and single-stream HT and VHT rates on a single radio chain and single antenna or antenna interface. On APs with external antennas, this feature uses the external antenna interface labeled **A0** or **ANT0** (radio chain 0); the other (one or two) antenna interfaces are left unused. If you are provisioning an 802.11n-capable AP, select the **Enable for Radio-0** or **Enable for Radio-1** check boxes in the **Single-Chain Mode** section to enable single-chain mode for the selected radio. AP radios in single-chain mode will transmit and receive data using only legacy rates and single-stream HT rates up to MCS 7. This feature is disabled by default.
9. (Optional) Define the AP name or SNMP location. The **AP list** section displays current information for an AP, and allows you to define additional parameters for your AP, such as AP Name, SNMP System Location.
10. Click **Submit** (Re-provisioning the AP causes it to automatically reboot).

Configuring a Remote AP

A Remote AP is recommended when the network between the AP and managed device is an un-trusted/non-routable network, such as the Internet. Furthermore, a Remote AP supports an internal DHCP server, while a Campus AP does not.

The following sections provide information on Remote authentication and RAP configuration:

- Configuring a Remote AP on page 644
- Configuring a Remote AP on page 644

Remote Authentication

The two most common ways to provision an AP for remote authentication are certificate-based AP provisioning and provisioning using a pre-shared key. Although both options allow for a simple secure setup of your remote network, you should make sure that the procedure you select is supported by your managed device, the AP model type and the end user's client software. If you must provision your APs using a pre-shared key, you need to know which managed device models you have that do not support certificate-based provisioning.

- **Certificate based authentication** allows a managed device to authenticate a AP using its certificates instead of a PSK. You can manually provision an individual AP with a full set of provisioning parameters, or simultaneously provision an entire group of APs by defining a provisioning profile which contains a smaller set of provisioning parameters that can be applied the entire AP group. When you manually provision an individual AP to use certificated-based authentication, you must connect that AP to the managed device before you can define its provisioning settings.
- Use **Pre-Shared Key (PSK) authentication** to provision an individual Remote AP or a group of Remote APs using an IKE PSK.

Remote AP Configuration

Following procedure describes how to configure Remote APs:

1. Select the managed device to which the AP will be provisioned.
2. Navigate to the **Configuration > Access Points** page.
3. Open the **Remote APs** tab
4. Select the new Remote AP from the **Remote AP** list, then click **Provision**.
5. In the **General** section, click the **AP Group** drop-down list and select the AP group to which this AP should be assigned. The AP group must have at least one virtual AP.
6. (Optional) Some AP models support an external antenna in addition to their internal antenna. If the AP you are provisioning supports an external antenna, the Provisioning window displays an additional **Antenna Parameters** section.
7. (Optional) To allow the Remote AP to use PEAP to authenticate to 802.1X networks, select **Show Advanced Options** under the **General** tab, then enter a user name and password in the 802.1X Parameter using PEAP section.

Verifying the Configuration

After the AP has been configured, navigate to **Dashboard** page and verify that the AP has an **up** status. The AP on your network *does not* appear in this table, it may have been classified as an inactive AP for any of the following reasons:

- The AP is configured with a missing or incorrect VLAN. (For example, the AP is configured to use a tunneled SSID of VLAN 2 but the controller does not have a VLAN 2.)
- The AP has an unknown AP group.
- The AP has a duplicate AP name.
- An AP with an external antenna is not provisioned with external antenna gain settings.
- Both radios on the AP are disabled.
- No virtual APs are defined on the AP.
- The AP has profile errors. For details, access the command-line interface and issue the command “show profile errors”.
- The GRE tunnel between the AP and the managed device was blocked by a firewall after the AP became active.
- The AP is temporarily down while it is upgrading its software. The AP will become active again after upgrading.

Validating and Optimizing AP Connectivity

The ArubaOS AP system profile contains multiple configuration settings to help you validate and optimize your AP connections to a managed device.

This section includes the following information:

- AP Health Checks on page 646
- Optimizing AP Connections over Low-Speed Links on page 646

AP Health Checks

The AP Health check feature uses ping probes to check reachability and latency levels for the connection between the AP and the managed device. The recorded latency information appears in the output of the **show ap ip health-check** command. If the managed device IP address becomes unreachable from the AP uplink, this feature records the time that the connection failed, and saves that information in a log file (tmp/ap_hcm_log) on the AP.



This feature is disabled by default, and is enabled by selecting the Health Check option in the AP system profile. For details see Optional AP Configuration Settings on page 708.

Optimizing AP Connections over Low-Speed Links

Depending on your deployment scenario, you may have APs or remote APs that connect to a managed device 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 managed device 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. Aruba APs locally process 802.11 probe-requests and probe-responses, but the 802.11 association process requires interaction with the managed device.

When deploying APs across low-speed or high-latency links, the following best practices are recommended:

- Connect APs and managed devices over a link with a capacity of 1 Mbps or greater.
- Maintain a minimum link speed of 64 Kbps per AP and per bridge-mode SSID. This is the minimum speed required for downloading software images.
- Adjust the bootstrap threshold to 30 if the network experiences packet loss. This makes the AP recover more slowly in the event of a failure, but it will be more tolerant to heartbeat packet loss.
- Prioritize AP heartbeats to prevent losing connectivity with the managed device.
- If possible, reduce the number of tunnel-mode SSIDs. Each SSID creates a tunnel to the managed device with its own tunnel keep alive traffic.
- If most of the data traffic will remain local to the site, deploy remote APs in bridging mode. For more information about remote APs, see Access Points on page 627.
- If high-latency links such as transoceanic or satellite links are used in the network, deploy a managed device 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.

The following topics provide information on configuring bootstrap threshold and prioritizing AP heartbeats:

- Configuring the Bootstrap Threshold on page 647
- Prioritizing AP Heartbeats on page 647

Configuring the Bootstrap Threshold

To configure the bootstrap threshold using the WebUI, enter a value into the **bootstrap threshold** field in the advanced AP system profile settings (For details, see Configuring the AP System Profile on page 718).

To configure this setting using the command-line interface, issue the command **ap system-profile <profile> bootstrap-threshold <bootstrap-threshold>**.

Prioritizing AP Heartbeats

To configure the AP heartbeat priority using the WebUI, enter a value greater than zero into the **Heartbeat DSCP** field in the advanced AP system profile settings (For details, see Configuring the AP System Profile on page 718).

To configure this setting using the command-line interface, issue the command **ap system-profile <profile> heartbeat-dscp <number>**.

AM Copy Optimization

Starting from ArubaOS 8.4.0.0, the AM Copy feature is significantly enhanced to reduce the burden on CPU and increase the AP performance. The optimization impacts the following features:

IDS Signature Match

This feature will no longer be reliable in matching the packet payload or sequence number. You need to disable the **wids-ampdu-optimization** parameter to detect the filtered packets and match all the packets with a given payload pattern or sequence number.

Frame Rate Anomaly Checks

The frame retry rate will be affected because the data received from the driver will now be filtered out. The frame retry rate will not be affected if **wids-ampdu-optimization** parameter is disabled.

AP Groups

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

Each AP is identified with an automatically-derived name. The default name depends on if the AP has been previously configured.

- The AP has not been configured—the name is the AP's Ethernet MAC address in colon-separated hexadecimal digits.
- Configured with a previous ArubaOS release—the name is in the format *building.floor.location*

You can assign a new name (up to 63 characters) to an AP; the new name must be unique within your network. For example, you can rename an AP to reflect its physical location within your network, such as "building3-lobby".

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 managed device are assigned. By using the "default" AP group, you can configure features that are applied globally to all APs.

Workflow for Configuring an AP Group

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.

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.



Reassigning an AP from an AP group requires a reboot of the AP for the new group assignment to take effect. Therefore, wait until there is little or no client traffic passing through the AP before reassigning it.

The tasks for configuring an AP group are:

1. Create an AP Group - You can create additional AP groups other than default.

   See Creating an AP group on page 648
2. Assign an AP to an AP Group - You can assign APs to that new group. A AP can belong to only one AP group at a time.

   See Assigning an AP to an AP Group on page 649
3. Assign channels to an AP Group - The country code in the AP Regulatory Domain profile determines supported channel and channel pairs for that specific AP. Any changes to the country code causes the valid channel lists to be reset to the defaults for the country.

   See Assigning Channels to an AP Group on page 649
4. Configure 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.

   See Configuring Channel Switch Announcement on page 650
5. Automatic Channel and Transmit Power Selection - Enable ARM to allow automatic channel and transmit power selection based on the radio environment.

   See Automatic Channel and Transmit Power Selection on page 650

Creating an AP group

The following procedure describes how to create an AP group:

1. In the **Managed Network** node hierarchy, select the managed device where the AP group are to be added.
2. Navigate to the **Configuration > AP Groups** menu.
3. Click **Add** below the AP Groups table.
4. In the **New AP Groups** window, enter the AP group name in the **New AP groups** field.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI command creates an AP group:

```
(host) [mynode](config) #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

Although you will assign an AP to an AP group when you first deploy the device, you can assign an AP to a different AP group at any time.



Once the **ap-regroup** command is executed, the AP automatically reboots. If the AP is powered off or otherwise not connected to the network or managed device, the executed command is queued until the AP is powered on or reconnected.

The following procedure describes how to assign a single AP to an existing AP group:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Access Points** menu.
2. Select the check box next to the AP and click **Provision**.
3. From the list of provisioning settings, click the **AP group** drop-down list and choose a new the AP group for the selected AP.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands assign a single AP to an existing AP group.

```
(host) [mynode](config) #ap-regroup {ap-name <name>|serial-num <number>|wired-mac
<macaddr>} <group>
```

Assigning Channels to an AP Group

The country code in the AP Regulatory Domain profile determines supported channel and channel pairs for that specific AP. Any changes to the country code causes the valid channel lists to be reset to the defaults for the country.

This section illustrates how to perform the following tasks for an AP group:

- Configure the “default” regulatory domain profile to use a valid country code. This will determine the available channels.
- Configure a 40 MHz channel (bonded pair) for the AP group’s 802.11a (5 Ghz) radio profile.
- Configure a 20 MHz channel for the AP group’s 802.11g (2.4 Ghz) radio profile.

The following procedure describes how to configure channels for an AP group:

1. In the **Managed Network** node hierarchy, select the managed device containing the AP group.
2. Navigate to the **Configuration** > **AP Groups** page.
3. Select the AP group to be configured.
4. Select the **Radio** tab from the AP group menu and click **Basic** accordion.
5. Click the **Radio mode** drop-down list and choose **ap-mode**.
6. Select 2.4 GHz or 5 GHz radio to be configured and click **Edit** in the **Valid Channels** field.
7. In the **Valid Channels** window, select the channels that will be supported by the AP group.
8. Click **OK**.

9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands configure channels for an AP group:

```
(host) [mynode](config) #ap regulatory-domain-profile default
   country-code US
(host) [mynode](config) #rf dot11a-radio-profile ht-corpnet-a
   channel 36+
(host) [mynode](config) #rf dot11g-radio-profile ht-corpnet-g
   channel 1
```



Country codes are generally specified in ISO 3166 format. To see what channels are available for a given country code, use the **show ap allowed-channels country-code <country-code>** command.

Configuring Channel Switch Announcement

When an AP changes its channel, an existing wireless clients may “time out” while waiting to receive a new beacon from the AP; the client must begin scanning to discover the new channel on which the AP is operating. If the disruption is long enough, the client may need to reassociate, reauthenticate, and request an IP address.

When CSA is enabled, the AP does not change to a new channel immediately. Instead, it sends a number of beacons (the default is 4) which contain the CSA announcement before it switches to the new channel. You can configure the number of announcements sent before the change.



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

The following procedure describes how to configure CSA:

1. In the **Managed Network** node hierarchy, select the managed device containing the AP group.
2. Navigate to the **Configuration** > **AP Groups** page.
3. Select the AP group to be configured.
4. Select **Radio** tab from the AP group menu and click **Advanced** accordion.
5. Select **Enabled** from the **CSA** drop-down list. This option can be enabled or disabled separately for 2.4 GHz and 5 GHz radios.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Automatic Channel and Transmit Power Selection

To allow automatic channel and transmit power selection based on the radio environment, enable 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 RF Planning and Channel Management on page 571.

AP Image Preload

The AP image preload feature minimizes the downtime required for a managed device upgrade by allowing the APs associated to that managed device to download the new images before the managed device actually starts running the new version.

This feature allows you to select the maximum number of APs that are allowed to preload the new software image at any one time, thereby reducing the possibility that the managed device may get overloaded or that network traffic may be impacted by all APs on the managed device attempting to download a new image at once.

APs can continue normal operation while they are downloading their new software version. When the download completes, the AP sends a message to the managed device, informing it that the AP has either successfully downloaded the new software version, or that the preload has failed for some reason. If the download fails, the AP will retry the download after a brief waiting period.

You can allow every AP on a managed device to preload a new software version, or also create a custom list of AP groups or individual APs that can use this feature. If a new AP associates to the managed device while the AP image download feature is active, the managed device will check that AP's name and group to see if it appears in the preload list. If an AP is on the list, (and does not already have the specified image in its Flash memory) that AP will start preloading its image.



Once a software version has been downloaded, another version cannot be downloaded until the AP reboots.

This section provides information on the following topics:

- Enable and Configure AP Image Preload on page 651
- View AP Preload Status on page 652

Enable and Configure AP Image Preload

Use the following procedures to enable and configure the AP Image Preload feature using the WebUI or CLI.

In the WebUI

1. To enable AP Image preload, navigate to **Maintenance > Software Management** in the **Managed Network** node hierarchy. Select the controllers that require upgrade from the **Controllers/Clusters** table.
2. To enable AP Image preload for a specific managed device, select the managed device and then navigate to **Maintenance > Software Management.**
3. Select **Preload software on access points** in **Installation Settings.**

In the CLI

To configure the AP image preload feature using the command-line interface, enter the following commands:

```
ap image-preload
    activate all-aps|specific-aps
    add {ap-group <ap-group> | ap-name <ap-name>}
    cancel
    clear-all
    delete {ap-group <ap-group> | ap-name <ap-name>}
    [partition <part-num>]
    [max-downloads <max-downloads>]
```

The parameters for this command are described in Table 118.

Table 118: *AP Image Preload Configuration Parameters*

Parameter	Description
activate	Issue the **ap image-preload activate** command to activate this feature, allowing APs in the preload list to start downloading their new image from the managed device.
all-aps	All APs will be allowed to pre download the image.
specific-aps	Only APs in the preload list will be allowed to preload the image.
add	Add individual APs or AP groups to the list of APs allowed to preload the image.
ap-group <group>	Add a group of APs to the preload list.
ap-name <name>	Add an individual AP to the preload list.
cancel	Cancel the AP preload and clear the preload list. Any APs downloading a new image at the time this command is issued will continue to download the file.
clear-all	Clear all APs from the preload list.
delete	Delete an individual AP or AP group from the preload list. **NOTE:** This command may be issued before or after preloading is activated. If it is executed after preloading has already been activated, any APs downloading a new image at the time this command is issued will continue to download the file. APs that are still waiting to preload will be removed from the preload list.
ap-group <group>	Remove the specified group of APs from the preload list
ap-name <name>	Remove an individual AP from the preload list
partition <partition-num>	Specify the partition from which the APs should download their images. By default, the APs will preload images from the managed device's default boot partition.
max-downloads <max-downloads>	Specify the maximum number of APs that can simultaneously download their image from the managed device. The default value is ten APs.

View AP Preload Status

You can monitor the current preload status of APs using the image preload feature using the **show ap image-preload status summary** command in the command-line interface. The output of this command contains the following information:

Table 119: *AP Image Preload Status Settings*

Column	Description
AP Image Preload State/Count	These two columns list the different possible preload states for APs eligible to preload a new software image, and the total number of APs in each state.

Column	Description
	■ **Preloaded**: Number of APs that have finished preloaded a new software image. ■ **Preloading**: Number of APs that are currently downloading the new image. ■ **Waiting**: Number of APs that are waiting to start preloading the new image from the managed device.
Count	This column lists the number of eligible APs currently in each preload state.
AP Name	Name of an AP eligible to preload a new software image.
AP Group	AP group of an AP eligible to preload a new software image.
AP IP	IP address of the AP.
AP Type	AP model type.
Preload State	Current preload state for the AP ■ **Preloaded**: The AP is finished preloading a new software image. ■ **Preloading**: The AP is currently downloading the new image. ■ **Waiting**: The AP is waiting to start preloading the new image from the managed device.
Start Time	Time the AP starting preloading an image.
End Time	Time the AP completed the image preload.
Failure Count	Number of times that the AP failed to preload the new image.
Failure Reason	In the event of an image preload failure, this column will display the reason that the image download failed.

AP Discovery Logic

In the earlier versions of ArubaOS, APs are predefined as either controller-based Campus APs or controller-less Instant APs. Each Campus AP is shipped with the ArubaOS manufacturing image and must connect to a controller in order to receive configurations. Campus APs can only run the ArubaOS image and cannot be converted into Instant APs. Each Instant AP is shipped with the Instant manufacturing image and must join an Instant AP cluster in order to receive configurations from a virtual controller. Instant APs run the Instant image and can also be converted into Campus APs.

Starting from ArubaOS 8.2.0.0, selected APs can run in both controller-based mode and controller-less mode. Based on the selected mode, the AP runs a different image:

- Controller-based APs run an ArubaOS image.
- Controller-less APs run an Instant image.

The following APs support both controller-based mode and controller-less mode:

- AP-203H
- AP-203R and AP-203RP
- AP-303H
- AP-365 and AP-367 access points

Each AP is shipped with a manufacturing image based on the Instant image, but containing reduced functions. When the AP is booted up with the manufacturing image, it enters the managed device and

Instant discovery process to determine if it will be upgraded to the controller-based mode (ArubaOS image) or controller-less mode (Instant image). After the managed device, Instant virtual controller, or Activate/AirWave/Central is discovered, the AP image is upgraded accordingly.

By default, controller discovery has a higher priority than Instant discovery. APs can discover the IP address of a managed device through one of the following methods. See Discovery of Controller on page 629 for more details on the different controller discovery options.

- Static controller discovery
- ADP
- DHCP server
- DNS server

Important Points to Remember

- APs can support up to 12 managed device IP addresses via DHCP/DNS discovery. APs attempt to connect to each managed device 10 times before switching to the next managed device.
- An AP can only be converted into a controller-based AP if the managed device to which it connects is running ArubaOS 8.2.0.0.
- If the AP cannot locate any managed device during the controller discovery process, it enters Instant discovery.

Preference Role

Users can predefine the AP mode by configuring the preference role. APs with the default preference role follow the standard discovery logic by attempting controller discovery before initiating Instant discovery. APs with the controller-less preference role bypass controller discovery and immediately initiate Instant discovery.

The following procedure describes how to set the AP preference role to controller-less:

1. Navigate to **Maintenance > Access Point > Convert to Instant Mode** in the WebUI.
2. Select the AP on which you want to set the preference role to controller-less.
3. Click **Convert to Instant Mode**.



This option is only available on Stand-alone controllers and managed devices.

You cannot convert a non-UAP model to an Instant AP. To convert a non-UAP model to an Instant AP, use the reset pin on the AP and reset the AP to factory default state.

The following CLI commands set the AP preference role to controller-less in the CLI:

```
(host) [mynode] #ap redeploy controller-less
   all
   ap-group
   ap-name
   ip-addr
   ip6-addr
   wired-mac
```



The **ap redeploy controller-less** command works only for UAPs and is applicable to AP-203H, AP-203R, AP-203RP, AP-303, AP-303H, 303P Series, AP-318, AP-344, AP-345, AP-365, AP-367, AP-374, AP-375, AP-377, AP-387, AP-534, and AP-535 access points only.

AP Deployment Policy

The AP deployment policy redirects the specified APs to the Instant discovery process, ensuring that the APs run only in controller-less mode. Users can predefine the AP deployment mode using the AP deployment policy.

The AP deployment policy can be configured on:

- APs in the specified IP address ranges—Policy is applied to the APs in the specified IPv4 or IPv6 address range. You can define up to 128 IPv4 and IPv6 address ranges for the AP deployment policy
- APs in the default AP group—Policy is applied to the APs in the default AP group.
- APs whose MAC address are included in the blacklist table—Policy is applied to the APs whose MAC addresses are included in the UAP blacklist table when the blacklist policy is enabled on the AP deploy profile.



You must enable the AP deploy profile to enforce the policies configured in the profile.

When the policy is enforced, the managed device automatically identifies the targeted AP, rejects the AP termination, and redirects the AP to upgrade to controller-less mode. The following CLI commands configure various AP deployment policies.

To enable the AP deploy profile, execute the following commands:

```
(host) [mynode] (config) #ap deploy-profile
(host) [mynode] (ap deploy-profile) #enable
```

To apply the AP deployment policy to the default AP group, execute the following commands:

```
(host) [mynode] (config) #ap deploy-profile
(host) [mynode] (ap deploy-profile) #default-ap-group
```

To apply the AP deployment policy to an IPv4 address range, execute the following commands:

```
(host) [mynode] (config) #ap deploy-profile
(host[mynode] (ap deploy-profile) #ip-range <start> <end>
```

To apply the AP deployment policy to an IPv6 address range, execute the following commands:

```
(host) [mynode] (config) #ap deploy-profile
(host) [mynode] (ap deploy-profile) #ipv6-range <start> <end>
```

To include AP MAC address to the UAP blacklist table, execute the following command:

```
(host) [mynode] (config) #uap-blacklist add mac-address <address> description
<description>
```

To apply the AP deployment policy to the blacklisted APs, execute the following commands:

```
(host) [mynode] (config) #ap deploy-profile
(host) [mynode] (ap deploy-profile) #blacklist
```

To remove the IP address range or default AP group from the profile, execute the following command:

```
(host) [mynode] (config) #no ap deploy-profile
```

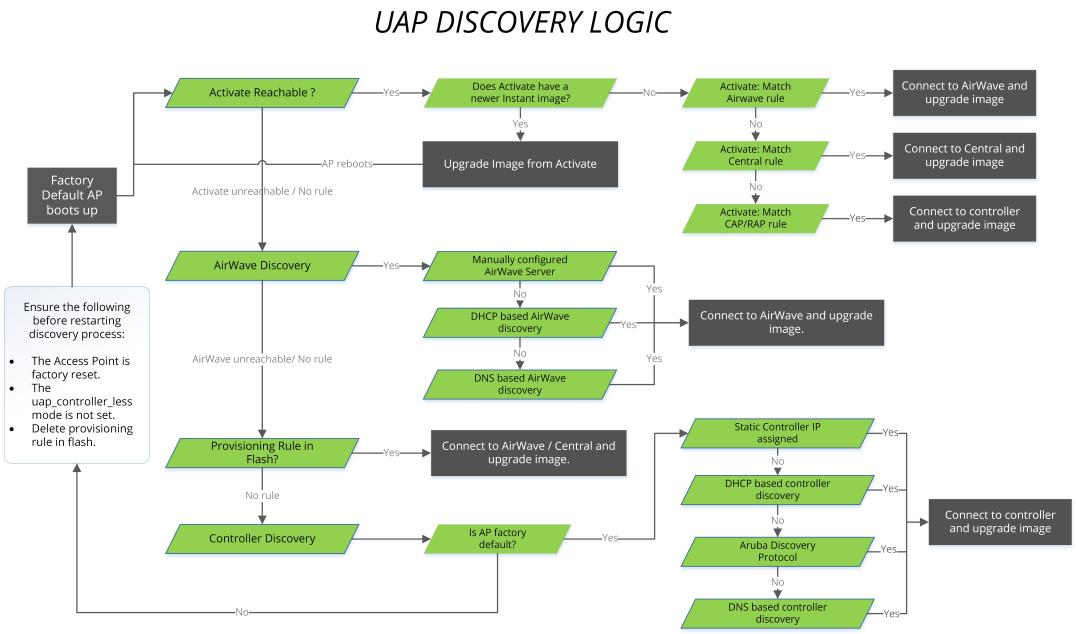
To view the complete list of IP address ranges to which the AP deployment policy is applied, execute the following command:

```
(host) [mynode] #show ap deploy-profile
```

Discovery Logic Workflow

The following steps describe the AP discovery logic:

Figure 85 *AP Discovery Logic*



1. When an AP boots up, it connects to Activate to obtain a provisioning rule.
2. If provisioning is already done by AirWave or Central, verify if a provisioning rule exists. If yes, the provisioning rule is saved in the flash memory. Compare the saved provisioning rule with the rule in Activate. If the rule in Activate is new, save the new provisioning rule in flash. For example, if the master and slave Instant APs obtain different AirWave addresses or if the master and slave Instant APs obtain a different AirWave or Central rule, the master Instant AP rule takes higher precedence.



Only the master Instant AP can apply provisioning rules to the Instant AP cluster.

3. If the rule is to perform a mandatory upgrade of the Instant AP, ensure to upgrade the Instant AP to the desired version. The master Instant AP executes the upgrade after a cluster is formed.
4. If the rule is to convert the Instant AP to Campus AP or Remote AP, the conversion takes effect for every Instant AP regardless of whether it is a master or a slave. This requires a manual registration of every master and slave Instant AP with Activate.
5. If there is no rule from Activate or if conversion to Campus AP or Remote AP fails, the master AP conducts local provisioning detection to check the local AirWave configuration.
   - If the AirWave server is configured and is in the configuration file, apply the server details. Otherwise, conduct a DHCP based AirWave or Central detection.
   - If DHCP-based AirWave is not found and the Instant AP is in factory default status, perform a DNS based AirWave discovery.
   - If none of the above methods can detect the AirWave server and if the Instant AP cannot connect to Activate, use the provisioning rule in flash.
6. If the AirWave or Central server is not found, or if the Instant AP is a slave, verify if the following conditions for local controller discovery are met:

- The Instant AP is factory reset.
- The **uap_controller_less** mode is not set.
- There is no provision rule saved in flash.

7. If the controller is found, the Instant AP sends a hello message to the controller and converts to a Campus AP.
8. When a master failover happens, the new master Instant AP connects to Activate to retrieve the provisioning rule. If the new master successfully obtains the provisioning rule, it applies this rule to the cluster.

Manual Upgrade

APs running in unprovisioned mode broadcast a special provisioning SSID to which users can connect to upgrade the AP manually. Upon connecting, users can access a local provisioning page in the WebUI to upgrade the AP to an ArubaOS or Instant image. See Controller-based AP using Manual Campus AP/Remote AP Conversion on page 660 and Controller-less AP using Manual Instant AP Conversion on page 662 for more details on upgrading APs manually.

Deployment Scenarios

This section describes various AP deployment scenarios in controller, Instant, remote, and hybrid networks.

See the following topics:

- Controller-based AP Deployments on page 657
- Controller-less AP Deployments on page 660

Controller-based AP Deployments

The following sections describe controller-based AP deployment scenarios.



Managed devices and APs are deployed in the same Layer 2 subnet.

Controller-based AP with AP Console Access

Users can deploy controller-based APs with console access, which allows them to modify the AP's provisioning settings through a direct console connection to the AP. This deployment scenario is typically used for troubleshooting in development/test networks and master controller assignment for static controller discovery. See Managing AP Console Settings on page 668 for more information on provisioning APs through a console connection.

To deploy a controller-based AP using an AP console connection:

1. Establish a console connection to the AP. See Managing AP Console Settings on page 668 for more details.
2. To access the AP console command prompt, press **Enter** when the AP displays the "Hit <Enter> to stop autoboot" message.
3. Enter the AP console password.
4. Execute one of the following APBoot commands to assign an IP address from which the AP can download the ArubaOS image:

   setenv serverip <ipaddr>: IP address of a TFTP server.

a. (Optional) To upgrade the image directly in partition **<n>** from **<file>**, execute the **os [<n>] <file>** command.

b. After the server IP address is assigned, enter **saveenv** to save your settings.

c. Reboot the AP using the **boot** or **tftpboot** command. The AP boots up with the ArubaOS image.

d. **setenv master <ipaddr>**: IP address of a managed device. This option is used for static controller discovery.

e. After the managed device is assigned, enter **saveenv** to save your settings.

f. Reboot the AP using the **boot** command. The AP boots up with the manufacturing image.

g. The AP enters the static controller discovery process.

h. If the assigned managed device is discovered, the AP connects to the managed device and downloads the ArubaOS image.

i. After the image is downloaded, the AP reboots.

j. The configuration synchronizes, and the AP runs in controller-based mode.

Controller-based AP in a Test Network

Users can provision controller-based APs in a test network before deploying the APs in a working network.



Managed devices in a test network can only be discovered using the ADP.

APs are upgraded to the ArubaOS image via ADP through the following steps:

1. The AP boots up with the manufacturing image in unprovisioned mode.
2. The AP enters the controller discovery process using ADP.
3. When a managed device is discovered, the AP connects to the managed device and downloads the ArubaOS image.

   If the AP cannot locate a managed device, it enters the Instant discovery process. To prevent the AP from upgrading to controller-less mode, you must make sure Instant virtual controllers, Activate, AirWave, and Central are not available to the AP. If the AP is not upgraded and there are no configuration changes with 15 minutes, the AP reboots and restarts the discovery process (step on page 658).
4. After the image is downloaded, the AP reboots.
5. The configuration synchronizes, and the AP runs in controller-based mode.

Controller-based AP in a New Controller-based Network

Users can deploy APs directly into a brand new controller-based network. APs are upgraded to the ArubaOS image using static/ADP/DHCP/DNS based controller discovery. See Discovery of Controller on page 629 for more details on the different controller discovery options.

APs are upgraded to the ArubaOS image via DHCP/DNS through the following steps:

1. The AP boots up with the manufacturing image in unprovisioned mode.
2. The AP enters the controller discovery process using DHCP/DNS.
3. When a managed device is discovered, the AP connects to the managed device and downloads the ArubaOS image.

APs attempt to connect to each managed device 10 times. If the AP fails to reach a managed device after 10 attempts, it reboots and restarts the discovery process (step on page 658).

4. After the image is downloaded, the AP reboots.
5. The configuration synchronizes, and the AP runs in controller-based mode.

See Controller-based AP in a Test Network on page 658 for details on ADP-based controller discovery.

Controller-based AP in an Existing Controller-based Network

Users can replace or add additional APs to existing controller-based networks. Newly deployed APs are be upgraded to the ArubaOS image using static/ADP/DHCP/DNS based controller discovery. See Discovery of Controller on page 629 for more details on the different controller discovery options.

See Controller-based AP in a Test Network on page 658 for details on ADP-based controller discovery. See Controller-based AP in a New Controller-based Network on page 658 for details on DHCP/DNS based controller discovery.

Controller-based AP in a Remote Deployment

Users can deploy controller-based APs in remote networks. APs in remote locations (Remote APs) connect to the Arubacontroller over the Internet using XAuth and IPsec. See Remote Access Points on page 844 for more information on Remote APs.

To deploy a controller-based AP in a remote site:

1. Login to the Mobility Master to add the AP to the managed device's Remote AP whitelist. See Managing AP Whitelists on page 68 for more details on adding APs to a Remote AP whitelist.
2. Place the AP in a remote site. The AP boots up with the manufacturing image in unprovisioned mode.
3. On your device, connect to the following provisioning SSID broadcasted by the unprovisioned AP:

   SetMeUp-xx:xx:xx
4. Open a web browser, and then navigate to the following URL:

   https://setmeup.arubanetworks.com
5. Under **Convert to**, select **RAP**.
6. Enter the IP address or host name of the managed device to which the Remote AP will be connected.
7. Click **Save**.

After the image is downloaded from the managed device, the AP reboots. The configuration synchronizes, and the AP becomes a Remote AP.

APs can also be converted into Remote APs using Aruba Activate. For more details, see Controller-based AP using Aruba Activate on page 659.

Controller-based AP using Aruba Activate

If the AP cannot locate any managed device during the controller discovery process, the AP enters Instant discovery. During the Instant discovery process, the AP attempts to connect through Activate if it cannot locate an Instant virtual controller. If Activate is provisioned to convert APs to controller-based Campus APs or Remote APs, any AP that connects to Activate is converted into a Campus AP or Remote AP. Refer to the latest *Aruba Activate User Guide* for details on configuring provisioning rules.

APs are upgraded to the ArubaOS image via Activate through the following steps:

1. The AP boots up with the manufacturing image in unprovisioned mode.
2. The AP enters the controller discovery process using static/DHCP/ADP/DNS based controller discovery.
3. If the AP cannot locate any managed device, it enters the Instant discovery process to locate an Instant virtual controller, Activate, AirWave, or Central.
4. The AP attempts to locate a virtual controller in an existing Instant AP cluster. If the AP cannot locate any virtual controllers, it attempts to connect through Activate.
5. If the AP connects to Activate, it checks for provisioning rules to convert into a Campus AP or Remote AP.



APs that connect to Activate are automatically upgraded from the manufacturing image to the latest ArubaOS image. Refer to the latest *Aruba Activate User Guide* for details on configuring provisioning rules.

   - If the AP converts into a Campus AP, it retrieves the IP address of the managed device. The AP connects to the managed device and downloads the ArubaOS image. After the image is downloaded, the AP reboots. The configuration syncs, and the AP becomes a Campus AP.
   - If the AP converts into a Remote AP, it retrieves the IP address of a managed device that has included the AP in its Remote AP whitelist. The AP connects to the managed device through an IPsec tunnel and downloads the ArubaOS image. After the image is downloaded, the AP reboots. The configuration synchronizes, and the AP becomes a Remote AP. For more information on Remote APs, see Remote Access Points on page 844.



   - The AP must be added to the managed device's Remote AP whitelist before it can retrieve the IP address of the managed device. For more details on adding APs to a Remote AP whitelist, see Managing AP Whitelists on page 68.

Controller-based AP using Manual Campus AP/Remote AP Conversion

If the AP cannot be converted into a Campus AP or Remote AP through Activate, users can connect to a special provisioning SSID broadcasted by the unprovisioned AP to manually convert the AP to a Campus AP or Remote AP through the WebUI. See Controller-based AP using Aruba Activate on page 659 for details on converting an AP into a Campus AP or Remote AP through Activate.

To manually convert an AP to a Campus AP or Remote AP in the WebUI:

1. On your device, connect to the following provisioning SSID broadcasted by the unprovisioned AP:
   SetMeUp-xx:xx:xx
2. Open a web browser. You will automatically be redirected to a special provisioning page in the WebUI to convert the AP.
3. Under Convert to, select **CAP** or **RAP**.
4. Enter the IP address or host name of the managed device to which the Remote AP or Campus AP will be connected.
5. Click **Save**.

After the AP is upgraded, it reboots as a Campus AP or Remote AP.

Controller-less AP Deployments

The following sections describe controller-less AP deployment scenarios.

Controller-less AP in an Instant Network

Users can deploy APs directly into a running Instant network, which is comprised of an Instant AP cluster and a virtual controller that manages the network. A virtual controller must be available before any AP can be upgraded through this deployment scenario. See **Customizing IAP Settings > Master Election and Virtual Controller** in the latest *Aruba Instant User Guide* for more details on electing a master in an Instant network.

APs are upgraded to the Instant image via a virtual controller through the following steps:

1. The AP boots up with the manufacturing image in unprovisioned mode.
2. The AP enters the controller discovery process using static/DHCP/ADP/DNS based controller discovery.
   - If the preference role is set to controller-less, the AP bypasses controller discovery and immediately enters Instant discovery (skip to step on page 661).
   - If a managed device is discovered, but the AP deployment policy is applied to this AP, the AP connects to the managed device and downloads the ArubaOS image. The managed device rejects the AP termination and redirects the AP to the Instant discovery process.
3. If the AP cannot locate any managed device, it enters the Instant discovery process to locate an Instant virtual controller, Activate, AirWave, or Central.
4. The AP attempts to discover a virtual controller in an existing Instant AP cluster.
5. If a virtual controller is discovered, the AP joins the existing Instant AP cluster and downloads the Instant image from the cluster.
6. After the image is downloaded, the AP reboots.
7. The configuration synchronizes, and the AP runs in controller-less mode.

Controller-less AP using Activate, AirWave, or Central

If the AP cannot locate a virtual controller in an existing Instant AP cluster, the AP attempts to connect to Activate, AirWave, or Central to upgrade the AP to the Instant image and form a new Instant AP cluster.



In this deployment scenario, Activate, AirWave, or Central must be accessible to the AP.

APs are upgraded to the Instant image via Activate, AirWave, or Central through the following steps:

1. The AP boots up with the manufacturing image in unprovisioned mode.
2. The AP enters the controller discovery process using static/DHCP/ADP/DNS based controller discovery.
   - If the preference role is set to **controller-less**, the AP bypasses controller discovery and immediately enters Instant discovery (skip to step on page 661).
   - If a managed device is discovered, but the AP deployment policy is applied to this AP, the AP connects to the managed device and downloads the ArubaOS image. The managed device rejects the AP termination and redirects the AP to the Instant discovery process.
3. If the AP cannot locate any managed device, it enters the Instant discovery process to locate an Instant virtual controller, Activate, AirWave, or Central.
4. The AP attempts to discover a virtual controller in an existing Instant AP cluster.
5. If the AP cannot locate a virtual controller in an existing Instant AP cluster, the AP attempts to locate Activate, AirWave, or Central to upgrade the image and form a new Instant AP cluster.



APs running the manufacturing image cannot form an Instant AP cluster.

If the AP locates Activate, it receives pre-configured provisioning rules to connect to AirWave or Central or convert into a Campus AP or Remote AP.



APs that connect to Activate are automatically upgraded from the manufacturing image to the latest Instant/ArubaOS image. Refer to the latest *Aruba Activate User Guide* for more details on configuring provisioning rules.

If the AP locates AirWave, it can be upgraded to the Instant image. If an enforced image upgrade rule is configured in AirWave, the AP is upgraded to the Instant image configured for the enforced upgrade rule. If no enforced upgrade rule is configured, the AP is upgraded to the latest Instant image in AirWave. After the AP is upgraded, it reboots in controller-less mode. Refer to the latest *AirWave 8.x User Guid*e for details on AP image upgrade.



All firmware must be uploaded to AirWave before the AP connects and downloads the Instant image. Refer to the latest *AirWave 8.x Aruba Instant Deployment Guide* for details on firmware upload.

If the AP locates Central, it can be upgraded to the Instant image through the **Maintenance > Firmware** page in the CentralUI. After the AP is upgraded, it reboots in controller-less mode. Refer to the latest *ArubaCentral User Guide* for more details on AP image upgrade.



Central syncs with Aruba Activate to retrieve the latest Instant image.

After the AP is upgraded to controller-less mode, it forms a new Instant AP cluster and converts into the master. Other un-deployed APs can join the cluster and upgrade to the Instant image.

Controller-less AP using Manual Instant AP Conversion

If the AP cannot be upgraded into an Instant AP through a virtual controller, Activate, AirWave, or Central, users can connect to a special provisioning SSID broadcasted by the unprovisioned AP to manually convert the AP to an Instant AP through the WebUI. See Controller-less AP in an Instant Network on page 660 and Controller-less AP using Activate, AirWave, or Central on page 661 for details on upgrading an AP to the Instant image using a virtual controller, Activate, AirWave, or Central.

To manually convert an AP to an Instant AP in the WebUI:

1. Login to your virtual controller.
2. Connect to the following provisioning SSID broadcasted by the unprovisioned AP:
3. **SetMeUp-xx:xx:xx**
4. Open a web browser. You will automatically be redirected to a special provisioning page in the WebUI to convert the AP.
5. Under **Access Point Setup**, select **Image File** or **Image URL** to upload the Instant image.
   - If you select **Image File**, click **Browse** to locate and select an Instant image file from your local file explorer.
   - If you select **Image URL**, enter the web address of the Instant image under **URL**.
6. Click **Save**.

After the AP is upgraded, it reboots in controller-less mode.

Troubleshooting the AP Discovery Logic

The following sections describe troubleshooting scenarios users may encounter in the AP discovery logic:

- Identifying the Controller Discovery Method on page 663
- The AP is Unable to Upgrade to the ArubaOS Image on page 663
- The AP is Unable to Upgrade to the Instant Image on page 665
- The SetMeUp Provisioning SSID is not Showing up on the Device on page 665
- The AP does not Reboot After an Upgrade Failure on page 665
- The Operational State of an AP Ethernet port goes down While Using a PoE Injector on page 665
- Accessing the CLI After an Image Upgrade on page 666

Identifying the Controller Discovery Method

APs can obtain the IP address of a managed device through one of the following methods:

- Static controller discovery
- ADP
- DHCP server
- DNS server

Execute the **show log provision** command on the AP to determine which controller discovery method was used to upgrade the AP to the ArubaOS image.

The AP is Unable to Upgrade to the ArubaOS Image

There are several reasons why an AP may not be able to upgrade to the ArubaOS image, even when a managed device is configured.

ADP is disabled

If the ADP is disabled on the managed device, the AP will not be able to locate any managed device on its own. Execute the **adp discovery enable** command in the CLI to enable ADP.

The AP preference role is set to controller-less

If the AP preference role is set to controller-less, the AP bypasses controller discovery and immediately initiates Instant discovery. Use one of the following methods to check if the AP preference role has been set to controller-less:

- Execute the **apboot> printenv** command on the AP console to view the current environment variable settings. The **uap_controller_less** field indicates if the preference role is set to controller-less:

- **uap_controller_less=1**: The controller-less preference role is enabled.
- **uap_controller_less=0**: The controller-less preference role is disabled.

See Managing AP Console Settings on page 668 for more details on APBoot commands.

- Check the boot up log from the AP console. If the preference role is set to controller-less, the *"ADP is disabled by uap_controller_less"* message appears.
- Execute the **show log provision** command on the AP console to view the AP provisioning logs. If the preference role is set to controller-less, the *"Controller discovery is disabled by ap-env uap_controller_less"* message appears.

The AP is not factory default

If the AP is not set to factory default (manufacturing image in unprovisioned mode), it cannot enter the controller discovery process. Use one of the following methods to check if the AP is factory default:

- Check the boot up log from the AP console. If the AP is not factory default, the *"Not factory_default ap. Do not run ADP."* message appears.
- Execute the **show log provision** command on the AP. If the AP is not factory default, the *"Controller discovery is disabled since UAP is not factory default status"* message appears.
- APs remain in unprovisioned mode after failing both controller and Instant discovery. If the AP is unable to upgrade to the ArubaOS or Instant image through the controller and Instant discovery process, it can be upgraded manually using the **SetMeUp-xx:xx:xx** provisioning SSID. Execute the **show network** command on the AP to check if the AP is connected to the **SetMeUp-xx:xx:xx** provisioning SSID. If the AP is connected to a different SSID, it is not factory default.

The managed device is not running the correct image version

An AP can only be converted into a controller-based AP if the managed device to which it connects is running ArubaOS 8.2.0.0. Managed devices that run a different version of ArubaOS do not support the AP discovery logic and cannot convert the AP to controller-based mode.

The AP is attempting to connect to a fake controller

If the AP fails to convert into a controller-based AP, the managed device to which it attempted to connect may be fake. Execute the **show log provision** command on the managed device to check if the AP failed to connect after 10 attempts.

FTP or TFTP permission is denied on the managed device

In order to download the ArubaOS image from the managed device, the AP must establish a FTP or TFTP connection to the managed device. If FTP or TFTP permission is denied on the managed device, the connection attempt is dropped, and the AP cannot download the ArubaOS image.

Use one of the following methods to check if FTP or TFTP permission is denied on the managed device:

- Execute the **show log upgrade** command on the AP. If FTP or TFTP permission is denied, the AP fails to connect to the managed device, and the following messages appear in the upgrade log:
  - **Connecting to <controller IP address>... failed: Connection time out.**
  - **Error: failed to retrieve image**
  - **Info: try with tftp to download the image.**
- Access controls can be applied to a managed device port to filter traffic between the managed device and the APs. Traffic must meet all criteria on the ACL in order to reach the managed device or AP. If it does not meet the criteria, the connection is dropped.
  - Execute the **show interface fastethernet <port> access-group** command to view the ACLs that have been applied to the port.
  - Execute the **show ip access-list <string>** command on the controller to view the detailed configuration for the ACL that has been applied to the port.
  - Execute the **interface fastethernet|gigabitethernet} <port> ip access-group <name> {in|out|session {vlan <vlanID>}}** command to apply an ACL to a port.
  - Execute the **no interface {fastethernet|gigabitethernet} <port> ip access-group <name> {in|out|session {vlan <vlanID>}}** command to remove an ACL from a port.
  - Execute the **ip access-list extended {<number>|<name>} deny <protocol>** command to reject traffic for a specific protocol. If you reject traffic for the UDP, TFTP traffic is dropped. If you reject traffic for the TCP, FTP traffic is dropped.

- Execute the **ip access-list extended {<number>|<name>} permit <protocol>** command to allow traffic for a specific protocol. If you allow traffic for the UDP, TFTP traffic can reach the managed device or AP. If you allow traffic for the TCP, FTP traffic can reach the managed device or AP.

The AP is Unable to Upgrade to the Instant Image

If the AP is marked as **CAP-only**, it cannot be upgraded to the Instant image. **CAP-only** APs can only be upgraded to the ArubaOS image.

Execute the **show log provision** command on the AP to check if your AP is **CAP-only**. If your AP is **CAP-only**, the **CAP-only sku** message appears.

The SetMeUp Provisioning SSID is not Showing up on the Device

The **SetMeUp-xx:xx:xx** provisioning SSID used for manual AP upgrade only appears on a device if an AP fails to upgrade to the ArubaOS or Instant image during the controller and Instant discovery process.

APs can support up to 12 managed device IP addresses for ArubaOS image upgrade. During the controller discovery process, the AP attempts to connect to each managed device 10 times until it reaches one successfully. Each connection attempt takes one minute. Depending on the number of managed devices that are located by the AP, it can take up to 120 minutes just to complete the controller discovery process.

Execute the **show log provision** command on the AP to track the progress of the controller discovery process. The provisioning log displays the total number of controller IPv4 and IPv6 addresses that have been detected by the AP, and the current stage of the discovery process.

The AP does not Reboot After an Upgrade Failure

If an AP fails to upgrade to the ArubaOS or Instant image during the controller and Instant discovery process, it enters a 15 minute reboot period. After the AP is rebooted, it restarts the discovery process. However, there are several conditions that can prevent the AP from completing the reboot:

- Keyboard input from the user.
- WebUI session connected to the AP (manual image upgrade).
- Pending image upgrade.
- Discovery of an AMP server.
- Discovery of a Central server.

Execute the **show log provision** command on the AP to determine why the AP has not rebooted. The provisioning log displays one of the following messages:

- **Could not reboot- upgrade is pending**
- **Could not reboot- keyboard input**
- **Could not reboot- airwave is found**
- **Could not reboot- UI session**
- **Could not reboot- central is found**

The Operational State of an AP Ethernet port goes down While Using a PoE Injector

Sometimes, while using a PoE injector, the output of the **show ap debug port status ap-name <ap-name>** command indicates that the operational state of an Ethernet port is down.

```
(host) [mynode] #show ap debug port status ap-name test-ap-225
```

```
AP "test-ap-225" Port Status
----------------------------
Port  MAC               Type  Forward Mode  Admin    Oper  Speed   Duplex  802.3az
PoE
----  ---               ----  ------------  -----    ----  -----   ------  -------
---
0     9c:1c:12:c0:ab:40 GE    N/A           enabled  up    1 Gb/s  full    disabled
N/A
1     9c:1c:12:c0:ab:41 GE    tunnel        enabled  down  N/A     N/A     N/A
N/A
STP  Portfast  TX-Packets  TX-Bytes  RX-Packets  RX-Bytes
---  --------  ----------  --------  ----------  --------
N/A  N/A       69707       37468577  107570      11707191
N/A  N/A       0           0         0           0
```

Execute the following command to verify if the AP is powered up using an 802.3af Power Sourcing Equipment.

```
(host) [mynode] #show ap debug system-status ap-name <ap-name> | include POE
Power Supply                 : POE-AF
```

The power supply, **POE-AF** indicates that a pre-standard PoE injector is being used to power up the AP. This causes the operational state of the Ethernet port to go down. To resolve this issue, enable the **ap2xx-prestandard-poe-detection** parameter from the respective provisioning profile of the AP:

```
(host) [mynode] (config) #ap provisioning-profile <profile-name>
(host) [mynode] (Provisioning profile "<profile-name>") #ap2xx-prestandard-poe-
detection
```



This parameter is applicable only for the 200 Series, 210 Series, 220 Series, 270 Series, or AP-203R access points and AP-203H, AP-205H, or AP-228.

Accessing the CLI After an Image Upgrade

After the AP image is upgraded, users cannot access the CLI for the first 180 seconds of uptime. The following message appears when a user attempts to login to the CLI during the initial 180 second uptime period:

```
login as: admin
System uptime is 147 seconds and CLI is not ready yet, please try again later.
```

AP Channel Scanning

The scanning algorithm is enhanced to reduce the delay between visits to some channel types, by changing their scan priority.

This section provides details on the following topics:

- Channel Types and Priority
- Scanning Optimizations
- Channel Group Scanning

Channel Types and Priority

A channel can belong to one or more channel types, depending on regulatory information and the activity that is detected on the channel. The frequency of visits to a channel depends on the priority of the channel type(s) to which it belongs. The following table describes the priority of channel types.

Table 120: *Channel Types and Priority*

Channel Priority	Channel Type	Description
One	DOS Channels	Channels where the AP is actively containing one more rogue devices in AM mode are marked with an **O** flag in the ARM CLI output (`show ap arm scan-times`).
Two	Active Channels	Channels where AP or Station activity has already been detected are marked with an **A** flag in the ARM CLI output and are visited in all scan-modes.
Three	Reg-Domain Channels	Channels that are in the AP's regulatory domain are marked with a **C** flag in the ARM CLI output and are visited in all scan modes.
Four	All Reg-Domain Channels	Channels that belong to any country's regulatory domain are marked with a **D** flag in the ARM CLI output and are visited only if the **scan-mode** is set to **All-Reg** or **Rare**.
Five	Unconventional Scan Channels	This new channel type category contains channels that belong to any country's regulatory domain, but with an unconventional scan direction. These channels are marked with a **J** or **M** flag in the ARM CLI output and are visited only if **scan-mode** is set to **All-Reg** or **Rare**.
Six	Rare Channels	Channels that do not belong to any country's regulatory domain are marked with a **Z** flag in the ARM CLI output. Rare channel scanning is done in the AM mode only if the rare scan mode is selected in the **AM Scanning** profile.

The country code in the AP Regulatory Domain profile determines supported channel and channel pairs for that specific AP. If there is a change in the country code, the valid channel list is reset to the default value for that country.



Use the **show ap arm scan-times ap-name <ap_name>** command to show scan state and flags for each channel.

Scanning Optimizations

The following optimizations enable the AP to achieve optimum RF monitoring. Unconventional Scans and Relative Priority of Channel Type Categories optimization apply to all AP types, but Channel Group Scanning optimization applies only to 200 Series models. All optimizations apply to AP and AM mode scanning.

This section provides information on the following topics:

Unconventional (direction) Scans

- Unconventional scans are 40 MHz scans of a channel in the direction away from the channel pair. For example, in the 44-48 channel pair:

- Conventional scans will be 44+ and 48-
- Unconventional scans will be 44- and 48+
- Unconventional scans are no longer interspersed with conventional scans. Unconventional scans operate with a lower frequency, because they belong to a new low priority channel type.

- Unconventional scans are performed in all-regulatory and rare scan modes. But these scans will not be performed if the scan mode is set to regulatory domain. This modification enables the AP to scan through active channels, regulatory channels, and all-regulatory channels faster.



Currently, 200 Series access points do not support unconventional or rare channel scanning.

Modifications in Scan Frequency

A modification is introduced to increase the frequency of visits to active and regulatory domain channels. Channel type categories are:

- DOS
- Active
- Regulatory domain
- All-regulatory domain
- Unconventional or rare



Unconventional or rare channels are merged for scanning.

Since 11ac AP radio can hear frames sub-channels when it performs an 80 MHz wide scan, scanning can be optimized by categorizing channels into scan groups, which are visited sequentially when a new primary channel is selected. This allows the AP scan through the list of channels faster, so that the delay between visits to channels in a group is reduced.

For more information on Channel Group Scanning, see Channel Group Scanning on page 668.

Channel Group Scanning

The following are the salient features of channel group scanning:

- Channel groups can be 80 MHz (4 channels), 40 MHz (2 channels), or 20 MHz wide (1 channel).
- Each channel is mapped to a group depending on the maximum width supported by that channel and the radio's capability. The maximum width supported by a channel is determined by the channel's membership in regulatory domain channel pairs or groups.
  - Channel 36, 40, 44, and 48 belong to 80MHz group
  - Channel 165 belongs to 20MHz group
- Channel groups are visited sequentially and the primary channel is rotated after each visit.
- Group scanning behavior is performed for 200 Series access points on A-band channels.



Scanning only once in each 80 MHz wide group allows the AP to scan through the channel list faster and also hear frames on sub-channels.

Managing AP Console Settings

An AP's provisioning parameters are unique to each AP. These parameters are initially configured on the Mobility Master and then pushed out to the AP and stored on the AP itself. Best practices are to configure an AP's provisioning settings using the Mobility Master WebUI. If you find it necessary to alter an AP's provisioning settings for troubleshooting purposes, you can do so using the WebUI and CLI, or alternatively, through a console connection to the AP itself.

To create a console connection to the AP:

1. Connect a local console to the serial port on the AP. You can connect the AP's serial port to a terminal or terminal server using an Ethernet cable, or connect the serial console port to a DB-9 adapter, then connect the adapter to a laptop using an RS-232 cable. For details on connecting to an AP's serial console port, refer to the installation guide included with the AP.
2. Establish a console communication to the AP, then power-cycle the AP to reboot it.
3. To access the AP console command prompt, press **Enter** when the AP displays the message "*Hit <Enter> to stop autoboot.*" If the autoboot countdown expires before you can interrupt it, turn the device off and then back on.
4. Once the AP boot prompt appears, enter the AP console password. You can issue any of the AP provisioning commands described in the Table 121. Remember, though these commands may be useful for troubleshooting, they are all optional and are *not* necessary for normal AP provisioning.

Table 121: *AP Boot Commands*



The list of AP boot commands may vary based on the APBoot image version.

Command	Description
boot	Boot the ArubaOS image from flash or USB, using currently saved environment variables. Any unsaved changes to the variables will be lost. This command has the following sub-parameters: ■ ap - Boot the ArubaOS image from flash. ■ usb:<path> - Boot the ArubaOS image from USB.
clear	Clear the ArubaOS image or other information. This command has the following sub-parameters: ■ all - Clear the cache and ArubaOS. ■ cache - Clear the cache sectors (mesh, Remote AP, Campus AP). ■ os <n> - Clear the image from the specified partition (default: 0). ■ prov - Clear provisioning image from the flash.
dhcp	Invoke DHCP client to obtain IP/boot parameters.
factory_reset	Reset the AP to factory default.
flash	Upgrade the boot image. **NOTE:** Exercise caution when using this command.
help	Help text for the AP boot commands.
mfginfo	Shows manufacturing information of the AP.
osinfo	Shows the ArubaOS image information on the AP.
ping	Check network connectivity.

Command	Description
printenv	List the environment variables and their current settings. AP boot environment variables are configured using the AP boot **setenv** command,
purgeenv	Reinstate AP boot configuration to factory default. This includes restoring the default environment variables.
reset	Perform RESET of the AP CPU.
saveenv	Save environment variables to persistent storage.
setenv ipaddr <ipaddr>	IP address to be assigned to the AP.
setenv netmask <netmaskip>	Netmask to be assigned to the AP.
setenv gatewayip <ipaddr>	IP address of the internet gateway used by the AP.
setenv name <ap name>	Name of the AP.
setenv group <group name>	Name of the AP group to which the AP should belong.
setenv master <ipaddr>	IP address of the AP's Mobility Controller.
setenv serverip <ipaddr>	IP address of the TFTP server from which the AP can download its boot image.
setenv dnsip <ipaddr>	IP address of the DNS server used by the AP.
setenv domainname <domain>	Domain name used by the AP.
tftpboot	Boot ArubaOS image over the network using TFTP protocol.
upgrade	Upgrade the APBoot or ArubaOS image. This command has the following sub-parameters: ■ boot <file> - Upgrade the APBoot image from <file>. ■ os [<n>] <file> - Upgrade the ArubaOS image in partition <n> from <file>. ■ prov - Upgrade provisioning image from <file>. **NOTE:** <file> can be a <TFTP-server-IP>:<path> or usb:<path>.
version	Displays the APBoot image version.

5. When you are finished, type **saveenv**and then press **Enter** to save your settings.



Other AP console commands may be available when accessing an AP directly through its console port, but these commands can cause configuration errors if used improperly and should only be issued under the direct supervision of Aruba technical support.

The example below configures an AP location and domain name using an AP console connection:

```
Hit <Enter> to stop autoboot: 0
```

```
apboot> <INTERRUPT>
apboot> setenv group corporate-2
apboot> setenv domainname mycompany.com
apboot> saveenv
apboot>boot
```

To view current AP settings using the AP console, issue the command **printenv <name>** where **<name>** is one of the variable names listed in Table 121, such as **ipaddr**, **dnsip** or **gatewayip**.

```
apboot> printenv domainname
domainname=mycompany.com
```

AP Console Password Protection

The ArubaOS AP console password feature helps protect systems that manage highly sensitive information, like financial and banking institutions, by requiring users to log in to the AP network with a password. The AP console password is enabled by default. Passwords must be 6 to 32 characters in length, and can include alphanumeric and special characters. If configured, you must enter this password to get AP console access. If not configured, the Mobility Master generates a default random password which can be viewed by executing the **encrypt disable** command followed by the **show ap system-profile <profile-name>** command.

The timeout feature is also supported as an added level of security. If there is no user input or activity during one timeout interval (default of 30 minutes), the user is logged out of the system. The timeout interval cannot be modified.

This section contains the following topics:

- Setting an AP Console Password on page 671
- Disabling Access to the AP Console on page 672

Setting an AP Console Password

The following procedure describes how to set a password in the WebUI:

1. In the Managed Network node hierarchy, navigate to **Configuration > System** page.
2. Select the **Profiles** tab.
3. Expand the **AP** profile in the **All Profiles** list, then select **AP System**.
4. Select the AP system profile you want to modify.
5. Open the **Advanced** tab, check the **Console Enable** check box.
6. In the **AP Console Password** field, enter the desired AP console password. Retype the password to confirm.
7. In the **Password for Backup** field, enter the password backup password for the console. Retype the password to confirm.
8. Click **Submit**.



Once the console is enabled, you do not need to enable it again. The console access is password protected.

The following CLI commands set the AP console password:

```
(host)[node] (config) #ap system-profile <profile>
(host)[node] (AP system-profile "<profile>") #console-enable
(host)[node] (AP system-profile "<profile>") #slow_timer_recovery
```

If the password is lost, and the AP is not connected to a managed device, the console can be reset using the reset button on the AP or the **factory_reset** AP boot command. If it is already connected to a managed device, the AP password can be changed under the **AP Console Password** field of the **AP System** profile in the WebUI, or using the **ap-console-password** parameter of the **ap system-profile** command in the CLI.

Disabling Access to the AP Console

Another way to protect your AP system is to completely disable access to the AP console under enabled mode.

The following procedure describes how to disable access to the console:

1. In the Managed Network node hierarchy, navigate to the **Configuration > System** page.
2. Select the **Profiles** tab.
3. Expand the **AP profile** in the **All Profiles** list, then select **AP System**.
4. Select the AP system profile you want to modify.
5. Open the **Advanced** accordion, deselect the **Console Enable** check box.
6. Click **Save**.

The following CLI commands disable access to the console:

```
(host)[node] (config) #ap system-profile default
(host)[node] (AP system profile "default") #no console-enable
```

Link Aggregation Support



Link aggregation is supported by 220 Series, 270 Series, 320 Series, 330 Series, 340 Series, AP-303P, 510 Series, 530 Series, and AP-555 access points.

All 220 Series, 270 Series, 320 Series, 330 Series, 340 Series, AP-303P, 510 Series, 530 Series, and AP-555 access points support link aggregation using either static port channel (configuration based) or Link Aggregation Control Protocol (protocol signaling based) link aggregation. These access points can optionally be deployed with LACP configuration to benefit from higher (greater than 1 Gbps) aggregate throughput capabilities. 330 Series, 340 Series, and 510 Series access points are limited to 1 Gbps on eth1 interface and hence eth0 interface cannot negotiate above 1 Gbps to form LACP.

The Mobility Master uses two different IP addresses for forwarding traffic to wireless clients associated to tunnel mode or decrypt-tunnel mode VAPs. One IP address is Mobility Master's IP address and the other is an unassigned IP address called GRE striping IP. Select the GRE striping IP address to ensure that a different physical interface is used by the load-balancing algorithm on the Ethernet switch. This enables the access points achieve greater than 1 Gbps throughput in both upstream and downstream directions.



AP LACP striping IP address need not be configured for APs terminating on a cluster.

On 200 Series and 270 Series access points, different IP addresses are used for different GRE tunnels between the AP and the LC. One LC IP address is used for tunnels corresponding to virtual APs using a 5 G radio and the other LC IP address is used for tunnels corresponding to virtual APs using a 2.4 G radio. By associating clients on both bands you can achieve more than 1 Gbps throughput.

A local AP LACP LMS map information profile that maps the device's LMS IP address to a GRE striping IP address. If the AP fails over to a standby or backup controller, the AP LACP LMS map information profile

on the new LC defines the striping IP address that the AP uses for link aggregation. This feature allows the access points to continue to support link aggregation to a backup controller in the event of a controller failover, even if the backup controller is in a different L3 network.

In previous releases, the GRE striping IP address was defined in the global AP system profile, which did not allow APs to maintain GRE striping tunnels if the AP failed over to a backup controller in a different L3 network.



If your topology includes a backup controller you must define GRE striping IP settings in the active and the backup controller For more information on LACP features in ArubaOS, see Configuring Port Channel LACP on page 119.

This section describes the following topics:

- Configuring LACP
- Important Points to Remember
- Troubleshooting Link Aggregation

Configuring LACP

To enable and configure LACP on 220 Series, 270 Series, 320 Series and 330 Series access points, specify the **LMS IP** address and configure the **GRE Striping IP** address in the AP LACP Striping profile. The **GRE Striping IP** value must be an IPv4 address owned by the Mobility Master that has the specified **LMS IP**. The **GRE Striping IP** does not belong to any physical or virtual interface on the Mobility Master, but the Mobility Master can transmit or receive packets using this IP.



The LMS IP address defined in the AP LACP profile or ap-lacp-striping command **must** be the same LMS IP address defined in the device's AP system profile. The LMS IP address in the device's AP system profile is used as a key to look up entries in the ap-lacp profile on the controllers to which an AP can connect.

The following procedure describes how to configure the LACP parameters in the AP System profile and AP LACP LMS map information profile.

1. In the Mobility Master node hierarchy, select the device.
2. Navigate to **Configuration > System**.
3. Select **Profiles** and expand the **AP** profiles menu.
4. Select the **AP LACP LMS map information** profile.
5. Select the **AP LACP Striping IP** check box to enable the **AP LACP striping IP** feature.
6. Enter a GRE striping IP address in the **IP** field. This IP address must be a in the device's subnet.
7. In the LMS field, enter the LMS IP address specified in the device's AP system profile in the **LMS** field. This LMS IP address *must* match the LMS IP address in Mobility Master's AP system profile.
8. Click **Pending Changes** and save your settings.
9. (Optional) Repeat these settings to configure LACP on a backup Mobility Master.

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System.**
2. Select **Profiles** and expand the **AP** profiles menu.

3. Select the **AP LACP LMS map information** profile.
4. Select the **AP LACP Striping IP** check box to enable the **AP LACP striping IP** feature.
5. Click **+**.
6. Enter a GRE striping IP address in the **IP** field. This IP address must be in the managed device's subnet.
7. In the **LMS** field, enter the LMS IP address specified in the device's AP system profile. This LMS IP address must exactly match the LMS IP address in the AP system profile configuration used by the device.
8. Click **OK**.
9. Click **Pending Changes** and save your settings.

When using high availability between two L3-connected controllers or two dual-mode HA controllers, you must define *two* different striping IPs (one in each controller subnet) to ensure that both the controllers will have striping IPs mapped to their corresponding LMS IP address.



When two controllers are both deployed in dual HA mode, each dual-mode controller acts as standby for the APs served by the other dual-mode controller. Each controller must therefore have two striping IPs, one for in each controller subnet. Two striping IP addresses are required for these topologies, even if the dual-HA controllers are located within the same subnet.

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System.**
2. Select **Profiles** and expand the **AP** profiles menu.
3. Select the **AP LACP LMS map information** profile.
4. Select the **AP LACP Striping IP** check box to enable the **AP LACP striping IP** feature.
5. Click +.
6. Enter a GRE striping IP address in the **IP** field. This IP address must be in the controller's subnet.
7. In the **LMS** field, enter the LMS IP address specified in the device's AP system profile. This LMS IP address must exactly match the LMS IP address in the AP system profile configuration used by the device.
8. Click **OK**.
9. Click +.
10. Enter a GRE striping IP address in the **IP** field. This IP address must be in the subnet of the other L3-connected or dual-mode HA controller.
11. In the **LMS** field, enter the LMS IP address specified in the device's AP system profile. This LMS IP address must exactly match the LMS IP address in the AP system profile configuration used by the device.
12. Click **OK**.
13. Click **Pending Changes** and save your settings.

The following CLI commands configure AP LACP and striping IP on a HA standby or backup LMS:

```
(host)[node] (config) #ap system-profile LACP
(host)[node] (AP system-profile "LACP") #lms-ip 192.0.2.1
(host)[node] (AP system-profile "LACP") #bkup-lms-ip 192.0.77.1
```

```
(host)[node] (AP system-profile "LACP") #exit
(host)[node] (config) #ap-lacp-striping-ip
(host)[node] (AP LACP LMS map information) #striping-ip 192.0.2.2 lms 192.0.2.1
(host)[node] (AP LACP LMS map information) #aplacp-enable
```

```
(bkup-host)[node] (config) #ap-lacp-striping-ip
(bkup-host)[node] (config) (AP LACP LMS map information) #striping-ip
192.0.2.16 lms 192.0.2.1
(bkup-host) [node] (config)(AP LACP LMS map information) #aplacp-enable
```

```
(bkup-host)[node] (config) #ap-lacp-striping-ip
(bkup-host)[node] (config) (AP LACP LMS map information) #striping-ip
10.1.1.14 lms 192.0.2.1
(bkup-host)[node] (config) (AP LACP LMS map information) #striping-ip
192.0.2.2 lms 192.0.2.1
(bkup-host) [node] (config)(AP LACP LMS map information) #aplacp-enable
```



If you are using High Availability between L3-connected or dual-mode controllers, you must configure **two** different striping IPs (one for each subnet) to ensure that both controllers will have striping IPs mapped to the corresponding LMS IP address.

Important Points to Remember

- In the upstream direction when the AP transmits GRE frames to the Mobility Master the bonding driver must be in active-active mode and not in the default active-standby mode to allow link aggregation.
- If an AP's uplink access switch ports are configured in static port-channel mode, then the AP will set the Ethernet bonding mode to static port-channel (xor mode) only if **gre-striping-ip** is configured. If **gre-striping-ip** is not configured, then the AP goes back to **active-standby** mode. In this scenario, the AP may go down depending on the behavior of the upstream switch.
- If an AP's uplink access switch ports are configured in dynamic LACP mode, the AP detects LACP-PDUs and automatically sets the Ethernet bonding mode to LACP. If **gre-striping-ip** is not configured, then the AP's Ethernet bonding mode will continue to be in LACP mode, but the AP will send GRE traffic only through one Ethernet port.
- In 320 Series and 330 Series access points, if AP uplink packet capture is taken, the downstream traffic will have sequence number in GRE header. Wireshark Aruba wlan decoder will not be able to decode these packets correctly since it looks for known Aruba GRE tunnel IDs.
- Ensure that the **gre-striping-ip** is unique and not used by any other host on the subnet.
- LACP support is limited to a use case where Enet 0 and Enet 1 ports of the AP are connected to a switch, and LACP is enabled on the two corresponding switch ports.
- The port priority is not applicable to the AP as both ports need to be used. This value is always set to the maximum numerical priority (0xFF), which is the lowest priority.
- The system priority is not configurable. It is set to the maximum numerical value (0xFFFF), which is the lowest priority. This leaves control of the aggregate to the upstream switch.
- The timeout value is not configurable.
- The key is not configurable and the default key value is 1.

- LACP cannot be enabled if wired AP functionality is enabled on the second port. You cannot enable LACP if the Enet 1 port is shutdown.

Troubleshooting Link Aggregation

The following show commands in the CLI can be used to troubleshoot Link Aggregation on 220 Series , 270 Series, 320 Series and 330 Series access points:

- **show ap debug lacp ap-name <ap-name>**—Using this command, you can view if LACP is active on an AP. It displays the number of GRE packets sent and received on the two Ethernet ports. Using this command with verbose option on 320 Series and 330 Series access points displays packet re-ordering statistics of each wlan client.
- **show ap database**—The output of this command includes an **LACP Striping** flags to indicate of the AP is configured with a LACP striping IP address,
- **show datapath tunnel**—Using this command on 220 Series/270 Series access points, you can verify if the 2.4 GHz tunnels are anchored on the **gre-striping-ip** (The GRE IDs for these tunnels are in a range between 0x8300 and 0x83F0) . On 320 Series and 330 Series access points, use the verbose option to verify that 5 Ghz tunnels have striping IP set in the column **StripIP** (The GRE IDs for these tunnels are in a range between 0x8200 and 0x82F0).
- **show datapath station**—On 320 Series and 330 Series access points, using this command displays the LACP sequence number sent in the GRE header of the last packet to the client. This information is displayed under **Seq** column.
- **show ap remote debug anul-sta-entries**—On 320 Series and 330 Series access points, using this command displays LAG enabled/disabled per station and data drops due to LAG packet reordering.
- **show datapath user**—Using this command, you can verify if the **gre-striping-ip** has an entry with the 'L' (local) flag.
- **show datapath route-cache**—Using this command, you can verify if the **gre-striping-ip** has an entry with the LC MAC.

2.4 GHz and 5 GHz Radio RF Management

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 ARM profile. If you would like ARM to dynamically select the best channel and transmission power for the radio, verify that the RF management profile references an active and enabled ARM profile. It can be useful to set the **Max Tx EIRP** parameter in the ARM profile to 127 (the maximum power level permissible) until it determines the signal-to-noise radio on the links. If ARM is active, the **Max Tx EIRP** can also be set to 127 to allow maximum power levels.

If you want to manually select a channel for each AP group, create separate 802.11a and 802.11g profiles for each AP group and assign a different transmission channel for each profile. For example, one AP group could have an 802.11a profile that uses channel 36 and an 802.11g profile that uses channel 11, and another AP group could have an 802.11a profile that uses channel 40 and an 802.11g profile that uses channel 9.

With the implementation of the high-throughput 802.11n standard, 40 Mhz channels were added in addition to the existing 20 MHz channel options. Available 20 Mhz and 40 Mhz channels are dependent on the country code entered in the regulatory domain profile. The newer VHT 802.11ac standard introduces 80 Mhz channel options.

Starting from ArubaOS 8.4.0.0, IEEE 802.11ax standard has been implemented that also supports 40 Mhz, 80 Mhz, and 160 Mhz channels in 5 Ghz frequency bands.



Changing the country code causes the valid channel lists to be reset to the defaults for the country.

Starting from ArubaOS 8.4.0.0, 5 GHz band includes 169 and 173 channels, for India only.

This section provides details on the following topics:

- Managing 2.4 GHz and 5 GHz Radio Settings on page 677
- Managing High Throughput Radio Settings on page 686
- RF Optimization on page 686
- RF Event Configuration on page 686

Managing 2.4 GHz and 5 GHz Radio Settings

This following sections explain the 2.4 GHz and 5 GHz radio settings and steps to configure the related parameters:

- Support for Dual 5 GHz Radio Mode on page 677
- Additional RF Management Settings in the WebUI on page 682
- Additional RF Management Settings in the CLI on page 685

Support for Dual 5 GHz Radio Mode

Starting from ArubaOS 8.3.0.0, 340 Series Series access points provide dual 5 GHz radio support. This feature is introduced for the AP-344 and AP-345 access points. In these AP models, by default, one radio operates in 5 GHz mode and the other in the 2.4 GHz mode.

The following points hold true when the **Dual 5GHz mode** option is enabled on 340 Series access points:

- Dual 5 GHz mode is a feature that you can enable at the AP level.
- Both radios operate in 5 GHz band and use the same dot11a-radio-profile.
- Both radios have the same Virtual APs that are allowed in the 5 GHz mode.
- The SSIDs of both radios are same.
- Both radios support 802.11a, 802.11n, 802.11ac 4x4:4SS 20/40/80/160 MHz (Max data rate of 1733 Mbps).
- The supported channels for the radios are within the prescribed Upper 5 GHz channel range (5.470–5.850 GHz) and Lower 5 GHz channel range (5.150–5.350 GHz).
- The middle 5 GHz channel range is not available for the radios.
- The APs support all forwarding modes—Tunnel, Decrypt Tunnel, Split, or Bridge—irrespective of the dual 5 GHz mode being enabled or disabled.
- AirMatch dynamically determines when the AP must operate in 5 GHz mode or dual band mode, when the **Dual 5GHz mode** option is set to **Automatic**. This feature is supported only on AP-345 access points.
- When **Dual 5GHz mode** option is enabled, the radios are not set to Air Monitor mode. The radios continue to operate in the AP mode.
- If ARM or AirMatch determines that due to limited channels, the AP (with **Dual 5GHz mode** enabled) must be forced into Air Monitor mode, one of the radios is forced to go into the Air Monitor mode.
- Depending on the Clarity client's radio band, the Clarity configuration is resent for the changed operation mode.

- If a Mesh AP operates in 2.4 GHz, you cannot set it to operate in the dual 5 GHz mode. This option is disabled for a Mesh AP.
- During configuration push for a Mesh AP, the **Dual 5GHz mode** option is set to disabled (override).
- MultiZone support is provided in ArubaOS 8.3.0.0 for the dual 5 GHz feature.

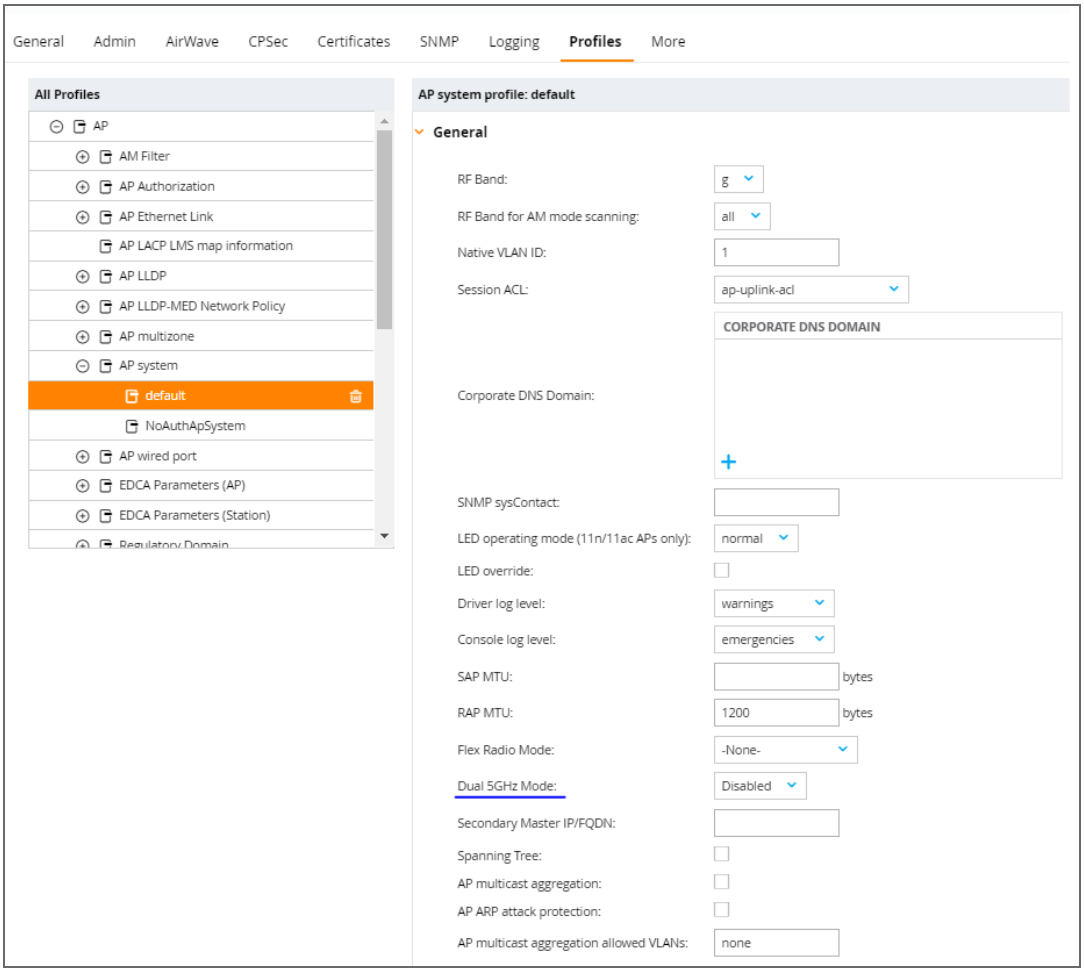
The new **Dual 5GHz mode** option helps to change one of the radios dynamically from a 2.4 GHz mode to a 5 GHz radio mode. You can set this feature to enable, disable, or automatic mode by using the WebUI or the CLI.

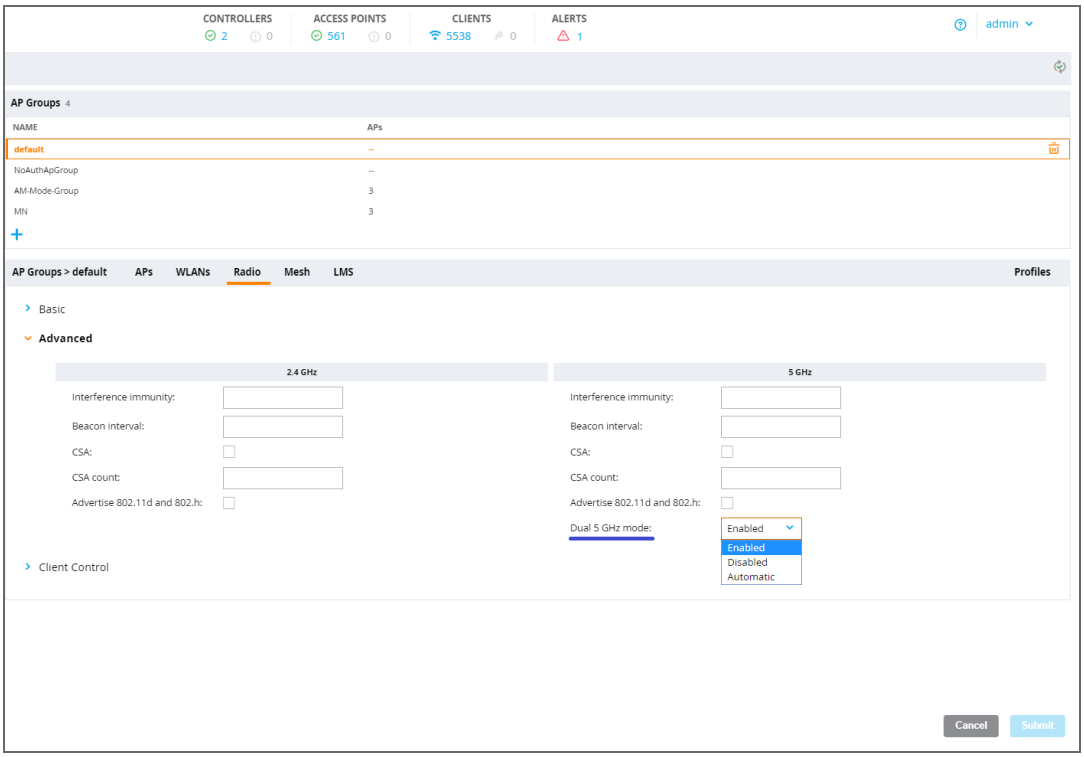
WebUI Enhancements to Support Dual 5 GHz Radio Mode

You can find the new **Dual 5GHz mode** option in the following paths in the **Managed Network** node hierarchy:

- Configuration > System > Profiles > AP > AP system profile.
- Configuration> AP Groups > Profiles > AP > AP system > AP system profile: <profile-name> (make sure the **Username > Preference > show advanced profiles** option is enabled)

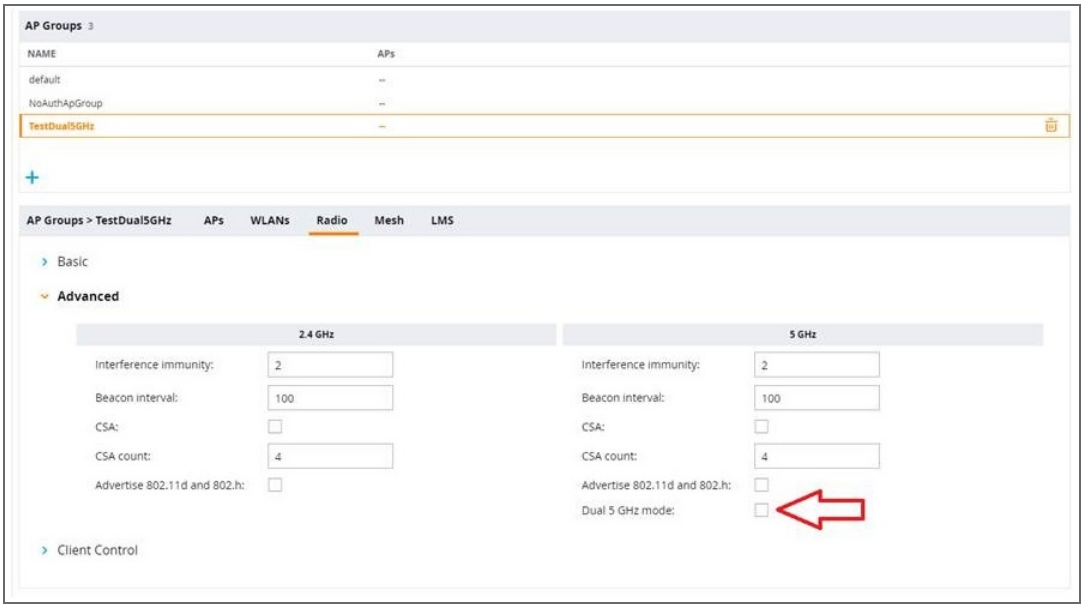
The representational ArubaOS 8.4.0.0 screenshots for these paths are as follows:





For a stand-alone controller or a master controller, the **Dual 5GHz mode** option is available in the **Configuration > AP Groups > <AP group profile-name> > Radio > Advanced** path.

The representational ArubaOS 8.3.0.0 screenshot illustrating the **Dual 5GHz mode** option for a stand-alone or a master controller is as follows:



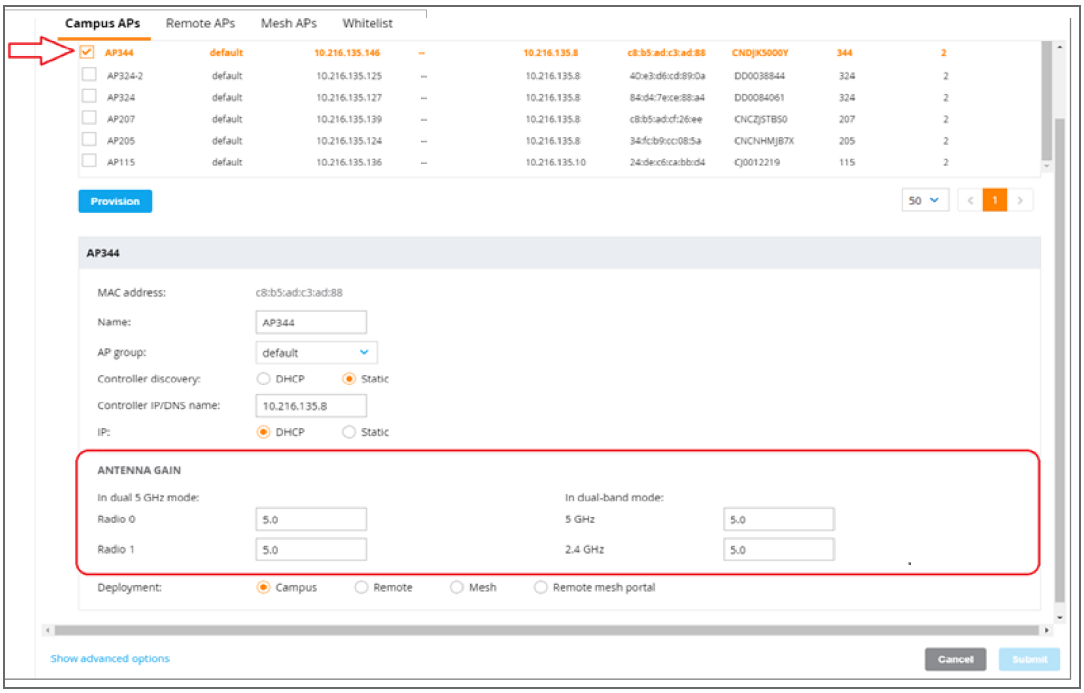
NOTE

The WebUI may change in a subsequent release for UI enhancement purposes.

When you select an AP-344 access point model in the Access Points table, you can see two additional **Antenna Gain** fields to set values for Radio 0 and Radio 1 for Dual 5 GHz mode. Find these parameters in the following paths:

- Configuration > Access Points > Campus APs.
- Configuration > Access Points > Remote APs.
- Configuration > Access Points > Mesh APs.

The representational ArubaOS 8.3.0.0 screenshot for the Antenna Gain parameter of Campus APs is as follows:



Use the following procedures to define and manage the most common 802.11a (5 GHz) and 802.11g (2.4 GHz) RF management settings using the WebUI or CLI:

- The following procedure describes how to manage the most common RF management settings:
- The following CLI commands set the Dual 5GHz mode option for 340 Series access points:

The following procedure describes how to manage the most common RF management settings:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > AP Groups** tab.
2. Select the name of an AP group from the AP groups tables.
3. Click the **Radio** tab below the AP groups tables to display the AP groups radio settings. The radio settings are divided into three sections, **Basic, Advanced**, and **Client Control.** The profile parameters in each section are described in Table 122.
4. Modify the desired settings, then click **Submit.**
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Table 122: *2.4 Ghz and 5 Ghz Radio Configuration Parameters*

Parameter	Description
Basic—Set the values for 2.4 GHz and 5 GHz radios	
Radio Mode	Access Point operating mode. Available options are: ▪ **am-mode**: Air Monitor mode

Parameter	Description
	■ **ap-mode**: Access Point mode ■ **spectrum-mode:** Spectrum Monitor mode The default settings is **ap-mode**.
Spectrum Monitoring	Select this option to convert APs using this radio profile to a hybrid APs that will continue to serve clients as an Access Point, but will also scan and analyze spectrum analysis data for a single radio channel. For more details on hybrid APs, see Spectrum Analysis on page 900. This option is available only when radio mode is **ap-mode**.
Transmit EIRP(dBm)	Select the maximum and minimum transmission power levels for the radio in 1 dBm increments. You can also set the maximum level to 127 for the regulatory maximum for that radio. Transmit power may be further limited by regulatory domain constraints and AP capabilities. This option is available only when radio mode is **ap-mode**.
Valid Channels	Click **Edit** to select a group of supported transmit channels for the 2.4 Ghz and 5 Ghz radios. The available channels depend on the regulatory domain (country). The available channels may be limited by the Channel Width setting. This option is available only when radio mode is **ap-mode**.
Scan mode	Air monitoring scan mode. Available options are: ■ **all-reg-domain**: Scan channels that belong to regulatory domain of any country ■ **rare**: Scan channels that do not belong to regulatory domain of any country ■ **reg-domain**: Scan channels that belong to regulatory domain of AP. The default settings is **all-reg-domain**. This option is available only when radio mode is **am-mode**.
Advanced—Set the values for 2.4 GHz and 5 GHz radios	
Interference Immunity	Set a value for non-Wi-Fi Interference Immunity. The default setting for this parameter is Level 2. When performance drops due to interference from non-802.11 interferes (such as DECT or Bluetooth devices), the level can be increased up to level 5 for improved performance. However, increasing the level makes the AP slightly "deaf" to its surroundings, causing the AP to lose a small amount of range. The levels for this parameter are: ■ Level 0: no ANI adaptation. ■ Level 1: noise immunity only. ■ Level 2: noise and spur immunity. ■ Level 3: level 2 and weak OFDM immunity. ■ Level 4: level 3 and FIR immunity. ■ Level 5: disable PHY reporting. **NOTE:** It is recommended not to use adjust the interference immunity feature without guidance from Aruba support.
Beacon Interval	Beacon Interval for the AP in ms. The supported range is 60-30000 ms, and the default value is 100 ms.

Parameter	Description
CSA	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. Enable this option to allow clients that support CSA to transition to the new channel with minimal downtime.
CSA Count	Number of channel switch announcements that must be sent prior to switching to a new channel. The default CSA count is 4 announcements. The allowed range is 1–16.
Advertise 802.11d and 802.11h	Enable the radio to advertise its 802.11d (Country Information) and 802.11h (Transmit Power Control) capabilities. This option is disabled by default.
Dual 5GHz mode	This option is available under 5 GHz radio settings. Select **Enable** if you want both radios to work in the 5 GHz mode. The available options are automatic, disabled, and enabled. **NOTE:** This option is applicable only to 340 Series access points (AP-344 access points and AP-345 access points).
Client Control—Set the values for 2.4 GHz and 5 GHz radios	
Client Match	Enable client match client bandsteering, load balancing, and enhanced AP reassignment for roaming mobile clients. For more information on this feature, see ClientMatch Overview on page 574

The following CLI commands set the Dual 5GHz mode option for 340 Series access points:

```
(host) [mynode] (config) #ap system-profile <profile-name>
(host) [mynode] (AP system profile "<profile-name>") #dual-5ghz-mode enabled
```

Use the following command to view the AirMatch reporting for an AP radio:

```
show ap debug airmatch reporting-radio
```

Use the following command to view the feasibility information about an AP debug:

```
show ap debug airmatch feasibility
```

Additional RF Management Settings in the WebUI

Additional Advanced RF management settings can be configured using the 2.4 GHz Radio Profiles and 5 GHz Radio Profiles.

1. In the **Mobility Master** node hierarchy, navigate to the **System> Profiles** tab.
2. Expand the **RF Management** profile menu.
3. Select either the **2.4 Ghz radio** profile menu, or the **5 Ghz radio** profile menu, then select the radio profile you wish to modify.
4. Modify the desired settings, then click **Submit.**
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Table 123: *5 GHz/2.4 GHz RF Management Configuration Parameters*

Parameter	Description
General 5 GHz/2.4 GHz Settings	

Parameter	Description
Radio Enable	Enable transmissions on this radio band.
Mode	Access Point operating mode. Available options are: ■ **am-mode**: Air Monitor mode ■ **ap-mode**: Access Point mode ■ **spectrum-mode:** Spectrum Monitor mode The default settings is **ap-mode**.
High throughput enable (Radio)	Enable/Disable high-throughput (802.11n) features on the radio. This option is enabled by default.
Very high throughput enable (Radio)	Enable/Disable very high-throughput (802.11ac) features on the radio. This option is enabled by default. **NOTE:** This parameter is only available in the 802.11a radio profile.
High efficiency enable (radio)	Enable/Disable High-Efficiency (802.11ax) features on the radio. This option is enabled by default.
Very high throughput rates enable (256-QAM)	Enable/Disable VHT rate on 2.4 GHz band providing 256-QAM modulation and encoding that allows for 600 Mbit/sec performance over 802.11n networks. **NOTE:** This parameter is only available in the 2.4 Ghz radio profile.
Channel	Transmit channel for this radio. The available channels depend on the regulatory domain (country). This parameter includes the following channel number configuration options for 20 MHz, 40 MHz and 80 MHz modes: ■ **20**: Select this option to disable 40 MHz mode and 80 Mhz mode and activate 20 MHz mode for the entered channel. ■ **40**: Entering a channel number and selecting the **40** radio button in the WebUI selects a primary and secondary channel for 40 MHz mode. When you use this option, the number entered becomes the primary channel and the secondary channel is determined by increasing the primary channel number by 4. For example, if you entered 157 into the **Channel** field and selected the **above** option, radios using that profile would select 157 as the primary channel and 161 as the secondary channel. ■ **80**; Entering a channel number and selecting the 80 Mhz radio button selects a primary and secondary channel for 80 MHz mode. If you select the spectrum monitoring checkbox on this profile page, the AP will operate as a hybrid AP and scan the selected channel for spectrum analysis data.
Non-Wi-Fi Interference Immunity	Set a value for non-Wi-Fi Interference Immunity. The default setting for this parameter is level 2. When performance drops due to interference from non-802.11 interferes (such as DECT or Bluetooth devices), the level can be increased up to level 5 for improved performance. However, increasing the level makes the AP slightly "deaf" to its surroundings, causing the AP to lose a small amount of range. The levels for this parameter are: ■ Level 0: no ANI adaptation. ■ Level 1: noise immunity only.

Parameter	Description
	■ Level 2: noise and spur immunity. ■ Level 3: level 2 and weak OFDM immunity. ■ Level 4: level 3 and FIR immunity. ■ Level 5: disable PHY reporting. **NOTE:** ■ Only 802.11n-capable APs simultaneously support both the Additional RF Management Settings in the WebUI on page 682 feature and a level-3 to level-5 Noise Immunity setting. Do not raise the noise immunity default setting on APs that do not support 802.11n unless you first disable the Channel Reuse feature. ■ It is recommended not to adjust interference immunity without guidance from Aruba support. ■ Only 300 Series access points (except AP-345) support this feature.
Spectrum Monitoring	Select this option to convert APs using this radio profile to a hybrid APs that will continue to serve clients as an Access Point, but will also scan and analyze spectrum analysis data for a single radio channel. For more details on hybrid APs, see Spectrum Analysis on page 900.
Max Channel Bandwidth	Sets the minimum channel bandwidth for APs associated to managed devices.
Min Channel Bandwidth	Sets the maximum channel bandwidth for APs associated to managed devices.
Min EIRP	The minimum transmission power level (in dBm) to be assigned to the AP radio(s).
Max EIRP	The maximum transmission power level 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.
eirp-offset	Manually adjust EIRP levels selected by the AirMatch algorithm by specifying a value from -6 to 6 dBm.
deploy-hour	Specify a number from 0-23 to select the hour during which AirMatch updates are sent to the APs (in 24-hour format). If the managed device to which the AP is associated is in a different time zone than Mobility Master, the AirMatch solution will be deployed according to the time zone of the managed device. If this parameter is set in both the AirMatch profile and the radio profile, the setting in the radio profile will take precedence.
Advanced 5 GHz/2.4 GHz Settings	
Transmit EIRP	Maximum transmit EIRP in dBm from 0 to 51 in .5 dBm increments, or 127 for regulatory maximum. Transmit power may be further limited by regulatory domain constraints and AP capabilities.
Enable CSA	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.

Parameter	Description
CSA Count	Number of channel switch announcements that must be sent prior to switching to a new channel. The default CSA count is 4 announcements.
Smart Antenna	Enable or disable the smart antenna feature on AP-335 access points.
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	The Spectrum Load Balancing feature helps optimize network resources by balancing clients across channels, regardless of whether the AP or the controller is responding to the wireless clients' probe requests. If enabled, the controller compares whether or not an AP has more clients than its neighboring APs on other channels. If an AP's client load is at or over a predetermined threshold as compared to its immediate neighbors, or if a neighboring Aruba AP on another channel does not have any clients, load balancing will be enabled on that AP. This feature is disabled by default.
Beacon Period	Beacon Period for the AP in msec. The range is 60-2000 msec, and the default value is 100 msec.
Beacon Regulate	Enable this setting to introduce randomness in the beacon generation so that multiple APs on the same channel do not send beacons at the same time.
ARM/WIDS Override	If selected, this option disables ARM and Wireless IDS functions and slightly increases packet processing performance. If a radio is configured to operate in Air Monitor mode, then the ARM/WIDS functions are always enabled, regardless of whether or not this check box is selected.
Management Frame Throttle Interval	Averaging interval for rate limiting management frames from this radio, in seconds. A management frame throttle interval of 0 seconds disables rate limiting.
Management Frame Throttle Limit	Maximum number of management frames that can come in from this radio in each throttle interval.
RX Sensitivity Threshold	RX sensitivity tuning based channel reuse threshold, in - dBm. If the Rx Sensitivity Tuning Based Channel reuse feature is set to static mode, this parameter manually sets the AP's Rx sensitivity threshold (in -dBm). The AP will filter out and ignore weak signals that are below the channel threshold signal strength. If the value for this parameter is set to zero, the feature will automatically determine an appropriate threshold.
Protection for 802.11b Clients	**(For 802.11g RF Management Profiles only)** Enable or disable protection for 802.11b clients. This parameter is enabled by default. Disabling this feature may improve performance if there are no 802.11b clients on the WLAN. WARNING: Disabling protection violates the 802.11 standard and may cause interoperability issues. If this feature is disabled on a WLAN with 802.11b clients, the 802.11b clients will not detect an 802.11g client talking and can potentially transmit at the same time, thus garbling both frames.

Additional RF Management Settings in the CLI

To view a complete list of 802.11a or 802.11g RF management profiles and their status:

```
[mynode]# show rf dot11a-radio-profile|dot11g-radio-profile
```

To view the settings of a specific RF management profile:

```
[mynode]# show rf dot11a-radio-profile|dot11g-radio-profile <profile-name>
```

Managing High Throughput Radio Settings

Each radio references a high-throughput profile that manages that AP's 40 Mhz tolerance settings. By default, a 5 GHz radio uses a high-throughput profile named **default-a** and a 2.4 GHz radio uses 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. For more information on configuring these settings, see High-Throughput APs on page 688.

RF Optimization

Each AP includes an RF Optimization profile that allows you to configure settings for detecting interference. The controller can detect interference near a wireless client station or AP is based on an increase in the frame retry rate or frame receive error rate.

The following procedure describes how to configure RF Optimization profiles:

1. Navigate to the **Configuration > System > Profiles** tab.
2. Select **RF Management** menu, then click **RF Optimization**.
3. Select the RF Optimization profile you want to edit or click **Add** and enter a name into the **Profile Name** dialog box to create a new profile.
4. Configure your RF Optimization radio settings then click **Submit**. Table 124 describes the parameters

Table 124: *RF Optimization Profile Parameters*

Parameter	Description
Station Handoff Assist	Allows the controller to force a client off an AP when the RSSI drops below a defined minimum threshold. Default: Disabled
RSSI Falloff Wait Time	Time, in seconds, to wait with decreasing RSSI before a de-authorization message is sent to the client. Maximum value: 8 seconds Default : 4 seconds
Low RSSI Threshold	Minimum RSSI above which de-authorization messages should never be sent. Default: 10
RSSI Check Frequency	Interval, in seconds, to sample RSSI. Default: 3 seconds

The following CLI command configures a RF Optimization profiles:

```
rf optimization-profile <profile>
```

RF Event Configuration

An AP's event threshold profile configures RSSI metrics, including high and low watermarks for frame error rates and frame retry rates. When certain RF parameters are exceeded, these events can signal excessive load on the network, excessive interference, or faulty equipment.



This profile and many of the detection parameters are disabled (value is 0) by default.

The following procedure describes how to configure RF event profiles:

1. Navigate to the **Configuration > Controller > Profiles** tab.
2. Select **RF Management** menu, then click **RF Event Thresholds**.
3. Select the RF Event Thresholds profile you want to edit or click **+** and enter a name into the **Profile Name** dialog box to create a new profile.
4. (Optional) Click **General** then select **Detect Frame Rate Anomalies** to enable or disables detection of frame rate anomalies. This feature is disabled by default.
5. (Optional) Click **Advanced** to configure the parameters described detailed in Table 125.
6. Click **Save**.

Table 125: *RF Event Thresholds Profile Parameters*

Parameter	Description
Bandwidth Rate High Watermark	If bandwidth in an AP exceeds this value, a bandwidth exceeded condition exists. The value represents the percentage of maximum for a given radio. (For 802.11b, the maximum bandwidth is 7 Mbps. For 802.11 a and g, the maximum is 30 Mbps.) The recommended value is 85%.
Bandwidth Rate Low Watermark	After a bandwidth exceeded condition exists, the condition persists until bandwidth drops below this value. The recommended value is 70%.
Frame Error Rate High Watermark	If the frame error rate (as a percentage of total frames in an AP) exceeds this value, a frame error rate exceeded condition exists. The recommended value is 16%.
Frame Error Rate Low Watermark	After a frame error rate exceeded condition exists, the condition persists until the frame error rate drops below this value. The recommended value is 8%.
Frame Fragmentation Rate High Watermark	If the frame fragmentation rate (as a percentage of total frames in an AP) exceeds this value, a frame fragmentation rate exceeded condition exists. The recommended value is 16%.
Frame Fragmentation Rate Low Watermark	After a frame fragmentation rate exceeded condition exists, the condition persists until the frame fragmentation rate drops below this value. The recommended value is 8%
Frame Low Speed Rate High Watermark	If the rate of low-speed frames (as a percentage of total frames in an AP) exceeds this value, a low-speed rate exceeded condition exists. This could indicate a coverage hole. The recommended value is 16%.
Frame Low Speed Rate Low Watermark	After a low-speed rate exceeded condition exists, the condition persists until the percentage of low-speed frames drops below this value. The recommended value is 8%.
Frame Non Unicast Rate High Watermark	If the non-unicast rate (as a percentage of total frames in an AP) exceeds this value, a non-unicast rate exceeded condition exists. This value depends upon the applications used on the network.
Frame Non Unicast Rate Low Watermark	After a non-unicast rate exceeded condition exists, the condition persists until the non-unicast rate drops below this value.

Parameter	Description
Frame Receive Error Rate High Watermark	If the frame receive error rate (as a percentage of total frames in an AP) exceeds this value, a frame receive error rate exceeded condition exists. The recommended value is 16%
Frame Receive Error Rate Low Watermark	After a frame receive error rate exceeded condition exists, the condition persists until the frame receive error rate drops below this value. The recommended value is 8%.
Frame Retry Rate High Watermark	If the frame retry rate (as a percentage of total frames in an AP) exceeds this value, a frame retry rate exceeded condition exists. The recommended value is 16%.
Frame Retry Rate Low Watermark	After a frame retry rate exceeded condition exists, the condition persists until the frame retry rate drops below this value. The recommended value is 8%.

The following CLI command configures an RF event profiles. The available parameters for this profile are detailed in Table 125.

```
rf event-thresholds-profile <profile>
```

High-Throughput APs

With the implementation of the IEEE 802.11ac standard, very-high-throughput can be configured to operate on the 5 GHz frequency band. High-throughput (802.11n) can be configured on both the 5 GHz and 2.4 GHz frequency bands. High-throughput is enabled by default, and can be enabled or disabled in the 802.11a and 802.11g radio profiles. For details, see 2.4 GHz and 5 GHz Radio RF Management on page 676.

Two different profiles advanced define settings specific to high-throughput APs, the **high-throughput radio** profile and the **high-throughput SSID** profile. Use the **High-throughput radio** profile to configure your APs to advertise intolerance of 40 Mhz operation (by default, this option is disabled, and 40 Mhz operation is allowed). This profile also allows you to enable the **CSD Override** feature. When you turn on CSD override, CSD is disabled and only one antenna transmits data, even if they are being sent to high-throughput stations. The **High-throughput SSID** profile configures the high-throughput SSID settings for 802.11n.

You must create and modify a high-throughput radio or high-throughput SSID profile to change default values for an AP radio, such as activating features not enabled by default, disabling features that are enabled by default, or modifying default values for configuration settings.



Stations are not allowed to use high-throughput with TKIP stand-alone encryption, although TKIP can be provided in mixed-mode BSSIDs that support high-throughput. High-throughput is disabled on a BSSID if the encryption mode is stand-alone TKIP or WEP.

This section describes the following topics:

- Configuring Advanced High-Throughput Radio Settings on page 688
- Configuring Advanced High-Throughput SSID settings on page 690

Configuring Advanced High-Throughput Radio Settings

Most deployments do not require manual configuration of the high-throughput radio profile.

However, you can configure advanced high-throughput radio profile settings using the WebUI or CLI.

In the WebUI

Perform the following task to configure advanced high throughput radio settings using the WebUI.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** window, expand the **RF Management** menu, then click **High-Throughput Radio**.
3. Select the High-Throughput Radio profile you want to edit or click + and enter a name into the **Profile Name** dialog box to create a new profile.
4. Configure the throughput settings described in Table 126.
5. Click **Save**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Table 126: *High-Throughput Radio Profile Configuration Parameters*

Parameter	Description
Honor 40MHz intolerance	When enabled, the radio will stop using the 40 MHz channels if the 40 MHz intolerance indication is received from another AP or station. Uncheck the **Honor 40 Mhz intolerance check box** to disable this feature. Default: Enabled
CSD override	Most transmissions to HT stations are sent through multiple antennas using cyclic shift diversity (CSD). When you enable the **CSD Override** parameter, CSD is disabled and only one antenna transmits data, even if they are being sent to high-throughput stations. This enables interoperability for legacy or high-throughput stations that cannot decode 802.11n CDD data. This option is disabled by default, and should only be enabled under the supervision of Aruba technical support. Use this feature to turn off antenna diversity when the AP must support legacy clients such as Cisco 7921g VoIP phones, or older 802.11g clients (e.g. Intel Centrino clients). Note, however, that enabling this feature can reduce overall throughput rates.
BSS Color	Enables different colors for each category of BSSIDs The Aruba 802.11ax based access points like AP-505, AP-515, AP-534. AP-535 and AP-555 support BSS coloring mechanism that helps identify the BSS from which a PLCP protocol data unit originates. Range : 0-63 and the default value is 0. **NOTE:** ■ 530 Series and 550 Series can detect and change the color automatically if the same color is detected for another BSS on the same channel. ■ The value of 0 means auto mode, that is, the AP will set the color by itself, finding any available color.
BSS Color Switch Count	The number of times the BSS color switch announcements are sent in beacons before switching to a new color. Range: 0-100 and the default value is 10.

In the CLI

Execute the following command to configure advanced high throughput radio settings:

```
(host)(config) #rf ht-radio-profile <profile>
(host)(config) #rf dot11a-radio-profile|dot11g-radio-profile <profile> high-
throughput-enable
```

Execute the following command to set the color of BSS and BSS Color Switch Count:

```
(host) [mynode] (High-throughput radio profile "default") #
bss-color
bss-color-switch-count
```

Execute the following command to view the set BSS color and BSS Color Switch Count:

```
(host) *[mynode] #show rf ht-radio-profile default-a
High-throughput radio profile "default-a" (Predefined (changed))
-------------------------------------------------------------------
Parameter                                       Value
---------                                       -----
40 MHz intolerance                              Disabled
Honor 40 MHz intolerance                        Enabled
CSD override                                    Disabled
VHT Bandwidth Signaling                         Disabled
VHT - Transmit Beamforming Sounding Interval    0 msec
BSS Color                       5
BSS Color Switch Count                10
```

Configuring Advanced High-Throughput SSID settings

Most deployments do not require manual configuration of the high-throughput SSID profile.

However, you can configure advanced high-throughput SSID profile settings or modify default SSID profile values using the WebUI or CLI.



De-A-MSDUs is supported with a maximum frame transmission size of 4 k bytes; however, this feature is always enabled and is not configurable. Aggregation is not currently supported.

In the WebUI

Perform the following task to configure advanced high-throughput SSID settings using the WebUI.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Expand the **Wireless LAN** menu.
3. Select **High-Throughput SSID**.
4. Select the High-Throughput SSID profile you want to edit or click + and enter a name into the **Profile Name** dialog box to create a new profile.
5. Configure the high-throughput SSID profile settings described in Table 127.

Table 127: *High-Throughput SSID Profile Parameters*

Parameter	Description
General	
High throughput enable (SSID)	Determines if this high-throughput SSID allows high-throughput (802.11n) stations to associate. Enabling high-throughput in a WLAN high-throughput SSID profile enables WMM base features for the associated SSID. Default: Enabled.
40 MHz channel usage	Enable or disable the use of 40 MHz channels. This parameter is enabled by default. Default: Enabled.
Very High throughput enable (SSID)	Enable or disable support for Very High Throughput (802.11ac) on the SSID. Default: Enabled.
80 MHz channel usage (VHT)	Enable or disable the use of 80 MHz channels on Very High Throughput APs. Default: Enabled.
Multi User Transmit Beamforming	
VHT - Multi User Transmit Beamforming	Enable or disable VHT Multi-User Transmit Beamforming. If this parameter is disabled, all other Multi-User Transmit Beamforming configuration parameters have no effect. Default: Enabled.
Transmit Beamforming	
VHT - Explicit Transmit Beamforming	Enable or disable VHT Explicit Transmit Beamforming for the 802.11ac-capable APs. When this parameter is enabled, the AP requests information about the MIMO channel and uses that information to transmit data over multiple transmit streams using a calculated steering matrix. The result is higher throughput due to improved signal at the beamforming (the receiving client). If this parameter is disabled, all other transmit beamforming settings will not take effect. Default: Enabled.
Advanced	
BA AMSDU Enable	Enable or disable Receive AMSDU in Block ACK (BA) negotiation. If enabled, AP denies clients from sending AMSDU using BA agreement. Default: Enabled.
Temporal Diversity Enable	When this setting is enabled and the client is not responding to 802.11 packets, the AP will launch two hardware retries; if the hardware retries are not successful then it attempts software retries. Default: Disabled.

Parameter	Description
Legacy stations	Control whether or not legacy (non-HT) stations are allowed to associate with this SSID. By default, legacy stations are allowed to associate. This setting has no effect on a BSS in which HT support is not available. Default: Enabled.
Low-density Parity Check	If enabled, the AP will advertise LDPC support. LDPC improves data transmission over radio channels with high levels of background noise. Default: Enabled.
MPDU Aggregation	Enable or disable MPDU aggregation. High-throughput APs are able to send A-MDPU, 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. Default: Enabled.
Max received A-MPDU size	Control the maximum size, in bytes, of an A-MPDU that can be received on this high-throughput SSID. Default: 65535 bytes.
Max transmitted A-MPDU size	Control the maximum size, in bytes, of an A-MPDU that can be sent on this high-throughput SSID. Range: 1576–65535 bytes.
Min MPDU start spacing	Minimum time between the start of adjacent MPDUs within an aggregate MPDU, in microseconds. Range: 0 (No restriction on MDPU start spacing), .25 µsec, .5 µsec, 1 µsec, 2 µsec, 4 µsec. Default: 0.
Short guard interval in 20 MHz mode	Enable or disable use of short (400 ns) guard interval in 20 MHz mode. 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: 400 ns (short) and 800 ns (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. Default: Enabled.
Short guard interval in 40 MHz mode	Enable or disable use of short guard interval (400 ns) in 40 MHz mode of operation. Default: Enabled.
Short guard interval in 80 MHz mode	Enable or disable use of short guard interval (400 ns) in 80 MHz mode of operation. Default: Enabled.

Parameter	Description
Supported MCS set	A list of MCS values or ranges of values to be supported on this SSID. The MCS you choose determines the channel width (20 MHz vs. 40 MHz vs. 80 MHz) and the number of spatial streams used by the mesh node. Range: 0–31. Default: 0–31. 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 MCS value of 16-23 are supported on RAP-155/11ac APs only. MCS value of 24-31 are supported on 320 Series APs only.
VHT - Supported MCS Map	Comma separated list of maximum supported MCS for spatial streams 1 through 4. Valid values for maximum MCS are 7, 8, 9, and '-' (if spatial stream is not supported). Maximum MCS of a spatial stream cannot be higher than the previous streams. If an MCS is not valid for a particular combination of bandwidth and number of spatial streams, it will not be used for Tx and Rx. Default: 9,9,9,9.
VHT - Transmit Beamforming Sounding Interval	Time interval in milliseconds between channel information updates between the AP and the beamformed client. Default: 25 msec. **NOTE:** This is applicable for 802.11ac-capable APs only.
Maximum VHT MPDU size	Maximum size of a VHT MPDU. Default: 11454 bytes.
Maximum number of MSDUs in an A-MSDU on best-effort AC	Set the maximum number of MSDUs in a TX A-MSDU on best effort AC. Default: 2. **NOTE:** In tunnel and decrypt-tunnel forwarding mode, TX A-MSDU is disabled if the value is set to 0. If the value is set to non-zero, TX A-MSDU is enabled and set to this value.
Maximum number of MSDUs in an A-MSDU on background AC	Set the maximum number of MSDUs in a TX A-MSDU on background AC. Default: 2. **NOTE:** TX A-MSDU is disabled if the value is set to 0. In decrypt-tunnel forwarding mode, TX A-MSDU on background AC is disabled and assigning any value has no effect.
Maximum number of MSDUs in an A-MSDU on video AC	Set the maximum number of MSDUs in a TX A-MSDU on video AC. Default: 2. **NOTE:** TX A-MSDU is disabled if the value is set to 0. In

Parameter	Description
	decrypt-tunnel forwarding mode, TX A-MSDU on video AC is disabled and assigning any value has no effect.
Maximum number of MSDUs in an A-MSDU on voice AC	Set the maximum number of MSDUs in a TX A-MSDU on voice AC. Default: 0. **NOTE:** TX A-MSDU is disabled if the value is set to 0. In decrypt-tunnel forwarding mode, TX A-MSDU on voice AC is disabled and assigning any value has no effect.

In the CLI

Execute the following command to configure advanced high-throughput SSID settings:

```
(host)(config) #wlan ht-ssid-profile <profile-name>
```

High-Efficiency (HE) APs

With the implementation of the IEEE 802.11ax standard, you can configure and improve spectrum efficiency and area throughput in dense deployment scenarios of APs or stations in both indoor and outdoor environments. You can configure High-Efficiency (HE) to operate on both the 2.4 GHz and 5 GHz frequency bands. HE is enabled by default, and can be enabled or disabled in the 802.11a, 802.11b, 802.11g, 802.11n, and 802.11ac radio profiles.

The 802.11ax certification requires the Wi-Fi Alliance's Agile Multiband Operation (MBO) certification. This certification enables Wi-Fi devices to better adapt to changing network conditions. It helps in improving resource utilization, balanced network load, and various other features.

The MBO certification also includes Cellular Data Capability (CDC). This provides APs or multi-mode capable STAs to advertise CDC support.

The **High-efficiency SSID** profile configures the High-Efficiency SSID settings for 802.11ax. Use the **High-efficiency SSID** profile to configure your APs to allocate the whole channel to a single user at a time or partition a channel to serve multiple users simultaneously.

This section describes the following topic:

Configuring Advanced High-Efficiency SSID settings

Most deployments do not require manual configuration of the High-Efficiency SSID profile. However, you can configure advanced High-Efficiency SSID profile settings or modify default SSID profile values using the WebUI or CLI.

In the WebUI

Perform the following task to configure advanced high efficiency SSID settings:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Expand the **Wireless LAN** accordion.

3. Select **High-Efficiency SSID**.
4. Select the High-Efficiency SSID profile you want to edit or click **+** and enter a name into the **Profile Name** dialog box to create a new profile.
5. Configure the High-Efficiency SSID profile settings described in Table 128.

Table 128: *High-Efficiency SSID Profile Parameters*

Parameter	Description
General	
High efficiency enable (SSID)	Determines if this High-Efficiency SSID allows High-Efficiency (802.11ax) stations to associate as HE-capable. Default: Enabled.
Advanced	
Dynamic fragmentation level	Control the level of Dynamic fragmentation level that is supported by the APs. Select one of the following options from the drop-down list: ■ Level-0: Does not support Dynamic Fragmentation ■ Level-1: Support for dynamic fragments that are contained within an S-MPDU. It does not support dynamic fragment within an A-MPDU that is not an S-MPDU. ■ Level-2: Support for dynamic fragments that are contained within an S-MPDU, and support for up to one dynamic fragment for each MSDU within an A-MPDU. ■ Level-3: Support for dynamic fragments that are contained within an S-MPDU, and support for up to four dynamic fragments for each MSDU within an A-MPDU. Default: Level-0 **NOTE:** This parameter is configurable only in bridge or D-Tunnel mode. **NOTE:** Configuring Level-2 and Level-3 options are targeted for future releases and are currently not supported in ArubaOS 8.5.0.0.
HE duration based RTS	Indicate the duration-based RTS value, in microseconds, in the HE capability. When the Transmission Opportunity (TXOP) is greater than the configured duration-based RTS value, RTS/CTS exchange value is used. The supported value range is 0 - 1023. Default: 1023
Individual TWT	Enable or disable individual TWT support. Default: Enabled.

Parameter	Description
HE TXBF	Enable or disable Transmit Beamforming (TxBF) in HE capability. Default: Enabled. **NOTE:** This parameter is targeted for future releases and is currently not supported in ArubaOS 8.5.0.0.
HE Supported MCS map	Comma separated list of maximum supported MCS for spatial streams 1 through 8. Valid values for maximum MCS are 7, 9, 11, and '-' (if spatial stream is not supported). Maximum MCS of a spatial stream cannot be higher than the previous streams. If an MCS is not valid for a particular combination of bandwidth and number of spatial streams, it is not used for Tx and Rx. Default: 11,11,11,11,11,11,11,11.
HE MU-OFDMA	Enable or disable MU-OFDMA in HE capability. Default: Enabled.
HE MU-MIMO	Enable or disable MU-MIMO in HE capability. Default: Enabled.
HE Guard Interval	Sets the supported GI. The supported ranges are 800 ns, 1600 ns and 3200 ns. Default: 800 ns.

In the CLI

Execute the following command to configure advanced High-Efficiency SSID settings:

```
(host)(config) #wlan he-ssid-profile <profile-name>
```

HE Pooling and Automatic Tri-Radio

ArubaOS supports High Efficiency (HE) dedicated radios, pooling of HE clients to HE-preferred radio, and automatic tri-radio mode.

AirMatch dedicates HE radios for ClientMatch to steer HE or 802.11ax capable clients to the dedicated radios. All 500 Series, 510 Series, 530 Series, 550 Series, 570 Series, and 570EX Series access points support HE pooling.

AirMatch supports automatic tri-radio mode, that is, two 5 GHz radios and one 2.4 GHz radio or the dual band mode of one 5 GHz radio and one 2.4 GHz radio on AP-555 access point.

Dedicated radios segregate HE clients and legacy clients to leverage 802.11ax efficiency with 802.11ax clients.

HE pooling requires all 802.11ax clients to operate on the same channel or sub-band of channels for realizing the 802.11ax performance gains. However, typical deployments have a mix of both 802.11ax capable APs (both dual and tri-radio) and legacy APs. In such a mixed deployment environment, HE pooling allows 802.11ax clients to associate with 802.11ax capable APs. Additionally, the 802.11ax clients associate to the particular radio on the AP which is operating on the selected channel or sub-band of channels.

The automatic tri-radio mode option for split 5 GHz tri-radio opmode decision is dynamically computed based on the network capabilities and capacity this avoiding manual intervention.

Configuring HE Pooling

The following procedure describes how to enable HE pooling:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles** tab.
2. In the **All Profiles** list, expand the AP menu, then select **AP system**.
3. Select the AP system profile you want to edit, or click **+** to create a new profile.
4. Configure the profile parameters described in AP System Profile Configuration on page 719.
5. Click **Advanced**.
6. Enable **HE Pooling**.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following commands enable HE pooling:

```
(host) [mynode] (config) #airmatch profile
(host) [mynode] (AirMatch profile) #he-pooling-client-density <he-pooling-client-
density>
```

The following commands disable HE pooling:

```
(host) [mynode] (config) #airmatch profile
(host) [mynode] (AirMatch profile) #he-pooling-client-density 0
```

The following command enables HE pooling on an AP or AP group. This is useful to include an AP or an AP group from participating in HE pooling:

```
(host) [md] (config) #ap system sample
(host) [md] (AP System profile "sample") #he-pooling-enable
```

The following command disables HE pooling on an AP or AP group. This is useful to exclude an AP or an AP group from participating in HE pooling:

```
(host) [md] (config) #ap system sample
(host) [md] (AP System profile "sample") #no he-pooling-enable
```

The following commands manage the signal thresholds for HE pooling steers. When a client moves from one radio to another, the RSSI at destination radio should not be significantly worse than its current association.

```
(host) [md] (config) #rf arm-profile sample-a
(host) [md] (Adaptive Radio Management (ARM) profile "sample-a") #cm-he-pooling-
signal-delta <cm-he-pooling-signal-delta>
```

The following commands manage the SNR for HE poling steers. When a client moves from one radio to another, the SNR at the destination radio should be as good as its current association:

```
(host) [md] (config) #rf arm-profile sample-a
(host) [md] (Adaptive Radio Management (ARM) profile "sample-a") #cm-he-pooling-snr-
thresh <cm-he-pooling-snr-thresh>
```

The following command shows the HE pooling statistics:

```
(host)(mynode) #show gsm lookup channel radio key radio_bssid 80:8d:b7:80:f4:b0 |
include he_pool
cfg_cm_he_pool_signal_delta :: 8
cfg_cm_he_pool_snr_thresh :: 30
he_pooling_moves :: 2
he_pooling_success :: 1
```

The following commands show the client-preference:

```
(host)(mynode) #show gsm lookup channel radio key radio_bssid 80:8d:b7:80:f4:d0 |
include client_pref
radio_client_preference :: 2
(host) (mynode) #show gsm lookup channel radio key radio_bssid 80:8d:b7:80:f4:b0 |
include client_pref
radio_client_preference :: 1
```

The following command shows the HE pooling-related steer statistics:

```
(host)(mynode) #show gsm lookup channel sta key sta_mac_address 6c:c7:ec:f2:b5:e2 |
include he_pool
he_pooling_moves :: 17
he_pooling_success :: 14
```

The following command shows the HE capability of a station in the client match state for the station:

```
(host) [mynode] #show ap arm client-match debug state client-mac 6c:c7:ec:f2:b5:e2
Mac :6c:c7:ec:f2:b5:e2
BSSID :80:8d:b7:80:f4:b0
ESSID :sathya-md-psk
HE Capable :Yes
AID/MU Client:3/1
dot11v/UCC active/Dualband/11v override:1/0/1/0
MBO: IE/Non pref chan cnt/CDC pres/CDC state:0/0/0/0
Client non pref channels :None specified
```

The following command shows the new steer type in the client match steer history:

```
(host) [mynode] #show ap arm client-match history
ARM Client match History
-------------------------
Time of Change Station Reason Status/Roam Time/Mode Signal(S/T/A/As) Band(S/T/A) Radio
Bssid(S/T/A) AP Name(S/T/A)
-------------- ------- ------ --------------------- ---------------- ----------- -----
------------- --------------
2020-01-10 02:50:01 6c:c7:ec:f2:b5:e2 HE-pooling Success/1/BTM-ACC -47/-47/-47/-39
5G/5G/5G 80:8d:b7:80:f4:b0/80:8d:b7:80:f4:d0/80:8d:b7:80:f4:d0 ap555/ap555/ap555
2020-01-10 02:45:00 6c:c7:ec:f2:b5:e2 HE-pooling No-Move/1/BTM-REJ1 -54/-54/-54/-40
5G/5G/5G 80:8d:b7:80:f4:b0/80:8d:b7:80:f4:d0/80:8d:b7:80:f4:b0 ap555/ap555/ap555
```

The following command shows the radio-client-preference as part of the VBR that is generated by Client Match:

```
(host) [mynode] #show ap virtual-beacon-report client-mac 6c:c7:ec:f2:b5:e2
Client MAC :6c:c7:ec:f2:b5:e2
HE Capable :Yes
Current association :ap555 (80:8d:b7:80:f4:d0)
-----------------
STA Beacon Report
-----------------
AP IP address Radio ESSID Signal (dBm) Last update Add time Channel/EIRP/Clients Flag
- ---------- ----- ----- ------------ ----------- -------- -------------------- ----
ap555 10.3.19.122 80:8d:b7:80:f4:d0 sathya-md-psk -50 Jan 10 03:38:13 Jan 10 02:40:49
161/18.0/1 *HD
ap555 10.3.19.122 80:8d:b7:80:f4:b0 sathya-md-psk -50 Jan 10 03:38:13 Jan 10 02:40:49
44/18.0/2 HN
ap344 10.3.19.114 c8:b5:ad:bb:7e:30 sathya-md-psk -72 Jan 10 03:38:04 Jan 10 02:42:37
157/15.0/0
ap344 10.3.19.114 c8:b5:ad:bb:7e:20 sathya-md-psk -72 Jan 10 03:38:04 Jan 10 02:42:37
44/15.0/0
ap555 10.3.19.122 80:8d:b7:80:f4:c0 sathya-md-psk -31 Jan 10 03:38:13 Jan 10 02:45:04
6/9.0/0 H
```

```
ap515 10.3.19.95 80:8d:b7:82:5a:10 sathya-md-psk -51 Jan 10 03:38:04 Jan 10 02:45:19
165/15.0/0 H
ap515 10.3.19.95 80:8d:b7:82:5a:00 sathya-md-psk -49 Jan 10 03:38:04 Jan 10 02:45:19
1/7.0/0 H
VBR Flags *- associated S-Stale U-unsupported channel H- HE radio N- non-HE preferred
radio D- HE preferred radio
```

The following command shows the configured HE-pooling thresholds and client preference in the radio state:

```
(host) [mm] (config) #show ap arm client-match debug state radio-bssid
80:8d:b7:82:5a:10
BSSID :80:8d:b7:82:5a:10
Name :ap515
cm/dot11v/blist_to/lb_client_thresh:1/1/3/30
steer_backoff/lb_thresh/lb_intvl/lb_client_thresh:300/20/5/30
lb_signal_delta/lb_snr_thresh/snr_thresh/max_steer_fails:5/30/10/2
channel/actual_eirp/num_vbr_nbr/num_clients/num_clients_recalc:165/15.0/1/1/1
Is Dummy/num_mu_clients/num_mu_clients_recalc/num_he_clients/Is MZ/Is HE:0/1/1/1/0/1
MU (SNR thresh/Client thresh):30/15
HE Pooling SNR Thresh/Signal Delta/Client Preference:30/8/all
```

The following command shows the VBR for a client present on an AP:

```
(host) #show ap arm virtual-beacon-report client-mac 98:09:cf:92:41:cd
Client MAC:98:09:cf:92:41:cd
Dual band:Yes
Active Voice:No
11v BTM capable:Yes
Steer capability:Steerable
Dual network capable:No
HE Capable:No
Current Association:c8:b5:ad:bb:7e:34
Virtual Beacon Report
---------------------
AP Channel Signal (dBm) EIRP Assoc HE Client-Preference
-- ------- ------ ----- ---- ----- -------------------
80:8d:b7:80:f4:b0 40 -39 15.0              Y    Non-HE
80:8d:b7:80:f4:d0 161 -39 15.0            Y    HE
80:8d:b7:82:5a:10 165 -42 15.0           Y    All
80:8d:b7:82:5a:00 1 -29 7.0                 Y    All
c8:b5:ad:bb:7e:20 36 -57 15.0       Y           All
c8:b5:ad:bb:7e:30 157 -57 15.0           Y    All
```

The following command shows the Client match summary with the overall statistics for the new steer algorithm:

```
(host) [mynode] #show ap arm client-match summary
SM: Sticky Moves, BM: Bandsteer Moves, LM: Load Balance Moves, MU: MUsteer Moves, VoM:
Voice Roam Moves, HM: HE Moves, HP: HE Pooling Moves
T: Total, S: Success
ACC: Accept, REJ#: Reject with reason #, TO: Timeout FA: False Accept
11v Move Format: (T/ACC/REJ1/REJ2/REJ3/REJ4/REJ5/REJ6/REJ7/REJ8/TO/FA)
Client Match Summary
---------------------
MAC SM (T/S) LM (T/S) BM (T/S) MU (T/S) VoM (T/S) HM (T/S) HP (T/S) Moves (T/S) Last
Move (Time/Rsn/Dur)) Device Type 11v Moves
- -------- -------- -------- -------- --------- -------- -------- ----------- --------
----------------- ----------- ---------
66:81:ca:a2:e9:bb 0/0 0/0 2/0 0/0 0/0 0/0 7/2 9/2 Jan 15 10:39:59 2020/HE-pooling/1
Linux 9/2/0/0/0/0/0/0/6/0/1/0
6c:c7:ec:f2:b5:e2 0/0 0/0 1/1 0/0 0/0 0/0 17/14 18/15 Jan 15 10:29:59 2020/HE-
pooling/X Android 7/7/0/0/0/0/0/0/0/0/0/0
Total clients:2
```

```
Sticky (T/S):0/0 Deauth (T/S):0/0 11v-BTM: 0/0/0/0/0/0/0/0/0/0/0/0
Load-balance (T/S):0/0 Deauth (T/S):0/0 11v-BTM: 0/0/0/0/0/0/0/0/0/0/0/0
Band-steer (T/S):3/1 Deauth (T/S):0/0 11v-BTM: 3/1/0/0/0/0/0/0/1/0/1/0
Voice-roam (T/S):0/0 Deauth (T/S):0/0 11v-BTM: 0/0/0/0/0/0/0/0/0/0/0/0
MU-Steer (T/S):0/0 Deauth (T/S):0/0 11v-BTM: 0/0/0/0/0/0/0/0/0/0/0/0
HE-steer (T/S):0/0 Deauth (T/S):0/0 11v-BTM: 0/0/0/0/0/0/0/0/0/0/0/0
HE-pooling (T/S):22/16 Deauth (T/S):12/10 11v-BTM: 13/8/0/0/0/0/0/0/5/0/0/0
```

The following procedure describes how to view client preference:

1. In the **Managed Network** node hierarchy, navigate to **Dashboard > Overview > Radios**.
2. Click the **Filter** and select **Client Preference**. The **Radios** table shows client preference.

The following CLI shows the per-radio statistic on clients and client-preference:

```
(host) [mm] #show ap arm client-match radio-summary
Radio Summary
--------------
Radio BSSID        AP Name           Phy Type         Client-Pref  Num Clients  Num HE
Clients
-----------        -------           --------         -----------  -----------  ------
--------
70:3a:0e:96:60:10  AP315-Airmonitor  5GHz (Non-HE)    Allow-All    1            0
70:3a:0e:96:60:00  AP315-Airmonitor  2.4GHz (Non-HE)  Allow-All    0            0
d0:15:a6:75:69:d0  AP505H            5GHz (HE)        HE-Pref      0            0
d0:15:a6:75:69:c0  AP505H            2.4GHz (HE)      Allow-All    0            0
f0:5c:19:1d:ff:c0  AP203R            5GHz (Non-HE)    Allow-All    0            0
f0:5c:19:1d:ff:d0  AP203R            2.4GHz (Non-HE)  Allow-All    0            0
90:4c:81:73:82:10  AP515-new-1       5GHz (HE)        Allow-All    4            3
90:4c:81:73:82:00  AP515-new-1       2.4GHz (HE)      Allow-All    0            0
80:8d:b7:80:f8:c0  AP555-new         5GHz (HE)        HE-Pref      8            7
80:8d:b7:80:f8:b0  AP555-new         2.4GHz (HE)      Allow-All    0            0
80:8d:b7:80:f8:a0  AP555-new         5GHz (HE)        Non-HE-Pref  1            0
80:8d:b7:80:b7:90  AP555-1           5GHz (HE)        Non-HE-Pref  0            0
80:8d:b7:80:b7:70  AP555-1           2.4GHz (HE)      Allow-All    0            0
20:a6:cd:34:bd:70  AP325-Airmonitor  5GHz (Non-HE)    Allow-All    0            0
20:a6:cd:34:bd:60  AP325-Airmonitor  2.4GHz (Non-HE)  Allow-All    0            0
90:4c:81:73:a7:20  AP515-new-2       2.4GHz (HE)      Allow-All    0            0
90:4c:81:73:a7:30  AP515-new-2       5GHz (HE)        Allow-All    1            0
c8:b5:ad:ba:f8:e0  AP345-new         2.4GHz (Non-HE)  Allow-All    0            0
c8:b5:ad:ba:f8:f0  AP345-new         5GHz (Non-HE)    Allow-All    0            0
Num Active Radios:19
Num HE-Pref Radios:2
Num Non-HE-Pref Radios:2
Num Clients:15
Num HE Clients:10
```

Loop Protection

The loop protect feature detects and avoids the formation of loops on the Ethernet ports of a Campus AP, Remote AP, or Mesh AP.

The loop protect feature can be enabled on all APs that have multiple Ethernet ports and it supports tunnel, split-tunnel, and bridge modes.

The loop protection feature prevents the formation of loops when:

- An unmanaged switch is connected to one port of an AP and a loop forms in the unmanaged switch.
- The WAN port (port 0) and either of ports 1, 2, 3, or 4, if it exists, in an AP are connected to the same switch.
- Multiple ports in an AP are connected to an unmanaged switch.

The loop protection feature transmits a proprietary loop detection packet on one Ethernet port of an AP at the configured loop-protect interval (default value is 2 seconds). The loop protect feature transmits the loop detection packet without a VLAN tag irrespective of whether the Ethernet port of the AP is connected in access mode or trunk mode. That is, for trunk mode, loop protect is supported only in the native VLAN.

- If the same packet is received on the same Ethernet port of the AP, a loop in the downstream switch is detected and the Ethernet port of the AP is shut down.
- If the same packet is received on the WAN port (port 0) of the AP, a loop between the Ethernet and WAN ports of the AP is detected and the Ethernet port of the AP is shut down.
- If the same packet is received on another Ethernet port of the AP, a loop between the Ethernet ports of the AP is detected and the Ethernet port of the AP port with lower priority is shut down. The Ethernet port with smaller port ID has high priority.

The Ethernet port of the AP that is shut down because of loop protection is marked with status **Loop-ERR**. A user can either the recover the shut down port from the managed device with manual intervention or enable automatic recovery mode and configure a automatic recovery interval. At the expiry of the automatic recovery interval, the **Loop-ERR** status of the Ethernet port is cleared and the Ethernet port is re-enabled automatically.

To prevent the downstream switch from dropping the loop detection packet, for example during broadcast storm state, if the AP takes longer time, or if the AP fails to detect a loop, a broadcast storm-control mechanism is provided as part of the loop protection feature. During broadcast-storm control, an AP counts the broadcast packets received on each of its Ethernet port and determines the packet rate in an interval. If the broadcast packet rate on one Ethernet port exceeds the configured threshold (default value is 2000 packets per second), the Ethernet port is shut down.

This section provides information on the following topics:

- Configuring Loop Protect on page 701
- Recovering Port on page 702

Configuring Loop Protect

The following procedure describes how to configure loop protect parameters in the AP wired port profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** page.
2. Select **AP > AP Wired Port**, and select the AP wired port profile you want to modify. The loop protect parameters for the profile are described in Table 129.

Table 129: *Loop Protect Parameters in AP Wired Port Profile*

Parameter	Description
Loop Protect Enable	Enables loop protection on AP wired ports.
Loop Detection Interval	Time, in seconds, to send loop detection packet. The supported range is 1 to 10 seconds and the default value is 2 seconds.
Storm Control Broadcast	Enables storm control broadcast. If the number of broadcast packets per second on one port in the AP exceeds the configured threshold, the port is shutdown.

Parameter	Description
Storm Control Broadcast Threshold	Storm control broadcast threshold in packets per second after which the port is shutdown. The default value is 2000 packets per second.
Auto Recovery Enable	Enables automatic recovery of the port in the AP. After the automatic recovery, if the loop re-occurs, then the port is shutdown again.
Auto Recovery Interval	Time, in seconds, to automatically recover the port in the AP. The supported range is 30 to 43200 seconds and the default value is 300 seconds.

The following CLI commands configure the loop protect parameters in the AP wired port profile:

```
(host) [mynode] (config) #ap wired-port-profile <profile>
(host) [mynode] (AP wired port profile "<profile>")#loop-protect-enable
(host) [mynode] (AP wired port profile "<profile>")#loop-detection-interval <loop-
detection-interval>
(host) [mynode] (AP wired port profile "<profile>")#auto-recovery-enable
(host) [mynode] (AP wired port profile "<profile>")#auto-recovery-interval <auto-
recovery-interval>
```

The following CLI commands display the status of the loop protect parameters in the AP wired port profile:

```
(host) [mynode] (config) #show ap wired-port-profile <profile>
AP wired port profile "<profile>"
-------------------------------
Loop Protect Enable                    Disabled
Loop Detection Interval                2
Storm Control Broadcast                Disabled
Storm Control Broadcast Threshold      2000
Auto Recovery Enable                   Disabled
Auto Recovery Interval                 300
```

Recovering Port

The following CLI commands manually recover a port of an AP in loop error state:

```
(host) [mynode] (config) #clear ap port ap-name <ap-name> <port>
```

Support for Port Bounce

Mobility Master provides support for the port bounce feature which enables a client to reinitiate a DHCP request when there is a VLAN change. This is achieved when a RADIUS server such as ClearPass Policy Manager sends Disconnect-Request with a Vendor Specific Attribute (VSA 40) to Mobility Master. Then, Mobility Master forwards the request to the device to trigger an interface shut down for a specified period. This allows the device to re-initiate a DHCP request for obtaining an IP address in the changed subnet.

The Disconnect-Request must include the following information:

- Calling Station-Id—MAC address of the user
- VSA—40
- Integer—0-60

VSA 40 represents **Aruba-Port-Bounce-Host**. The integer value indicates the time in seconds for which Mobility Master must shut the interface down. If the integer value received is greater than 60, then the port is shut down for default value of 12 seconds. If the value is 0, then the port is not shut down.



During a port bounce, the client connected to the interface is removed from the user table and is added back after the port is up.

Execute the following command to view the security logs during and after a port bounce:

```
[mynode] #show log security all | include bounce
```

The following sample shows the output during a port bounce:

```
Sep 14 22:22:46  authmgr[539]: <124004> <DBUG> |authmgr|  Sending port bounce request
for User mac 34:e6:d7:24:c8:3b
Sep 14 22:23:22  authmgr[539]: <124004> <DBUG> |authmgr|  Port Bounce succeeded for
User Mac 34:e6:d7:24:c8:3b
```

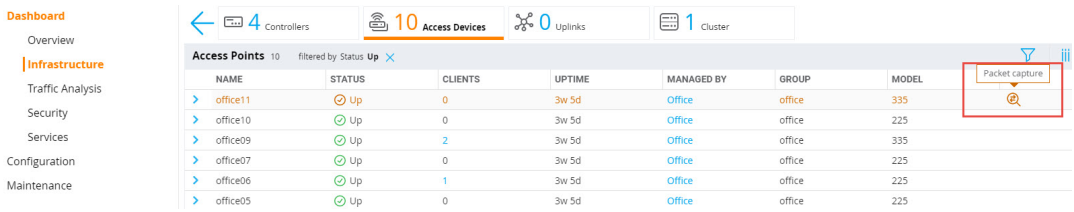
AP Packet Capture

Starting from ArubaOS 8.4.0.0, this feature allows you to manually start and stop capturing 802.11 Wi-Fi packets sent and received by an access point. This feature is supported only on an access point that is Up.

The following procedure describes how to capture packets:

1. In the **Managed Network** hierarchy, navigate to **Dashboard >Infrastructure > Access Devices.**
2. In the **Access Points** table, hover over an access point and then click on packet capture icon of an access point. The **Packet Capture** dialog box is displayed.

Figure 86 *AP Packet Capture*



3. Enter the **Target IP address** and the **Port** number to which you want to send the captured packets to.
4. Select a format from the **Format** drop-down list. The default format is **pcap**.
5. Select the bandwidth using the **Band** drop-down list. You can either select **2.4 GHz** or **5 GHz**. The default value is **2.4 GHz**.
6. Click **Start** to start capturing the packets.

Capturing packets can be stopped or paused using the options, **Pause** or **Stop** from the **Packet Capture** dialog box.



In Remote APs, the captured packets are sent from the AP to the controller. Then, the controller routes the packets to the target IP address. Hence, issue the **ap packet-capture open-port <port>** command to allow access to the UDP port to capture packets and then issue the **ap packet-capture close-port <port>** after capturing packets.

Green AP

Green AP is a feature that helps save energy consumption from common equipment in various areas like airports, offices, universities, hotels and so on. Based on the feeds, the Green AP feature dynamically enables, disables, or reduces functionality of an allocated AP to reduce the consumption of energy.

NetInsight provides the feed for ArubaOS to move the APs into deep-sleep mode or to wake up the APs from deep sleep mode. ArubaOS is responsible for maintaining the state of the APs and forwarding the AMON telemetry in different state.

In a Campus AP setup, the Mobility Master will communicate between NetInsight and the APs. For example, if NetInsight determines that a list of APs are to be put in deep-sleep or power saving mode, NetInsight sends the list to Mobility Master and then, the Mobility Master forwards the request to the APs through the managed device. The AP then, decides to either accept or reject the deep sleep request and sends the status back through AMON messages to the Mobility Master. This is again communicated to NetInsight.

The APs will not fall into deep-sleep mode in the following scenarios:

- The AP does not support WoL functionality.
- In MultiZone where APs need to provide wireless services for Datazone.
- The AP is preloading image.
- The AP is writing flash
- The APs have pending STAs.
- Wired AP is enabled on an AP.
- AP is 802.1X enabled.

Before the AP falls into deep-sleep mode, it performs the following actions:

- Bring down all the virtual APs.
- Send a warning syslog message
- Remove all connections to managed devices.
- Set the reboot reason. This is set to ensure that when the AP wakes up from the deep-sleep mode, this reboot reason indicates that the AP has recovered from deep-sleep mode.
- AP falls into deep-sleep mode.

Whenever the AP wakes up from the deep-sleep mode, the AP gets rebooted and the reason for the reboot is logged as, AP is waken up from deep-sleep mode.

AP-514, AP-515, AP-534, AP-535, AP-555, AP-504, AP-505, AP-505H, AP-518, and 570 Series access points support the Green AP, a power saving feature.



APs wake up automatically every 2 hours using the BLE process and report the status to NetInsight and if the NetInsight communicates that they need to be put back into deep-sleep mode, then the APs are again put into deep-sleep mode.

Limitations

Green AP feature is not supported for the following:

- Legacy APs without WoL support
- Instant APs

- Remote APs
- Mesh portals and mesh points
- Stand-alone controllers deployment

Configuring Green AP

The following CLI commands configure the green AP. These CLI commands are not available on the managed device and stand-alone controller deployment.

The following CLI commands put the AP in deep sleep mode:

```
(host) [mynode] #ap deep-sleep
ap-name <ap-name>
ip-addr <ip-addr>
ip6-addr <ip6-addr>
mac-list <mac_list>
wired-mac <wired_mac>
```

The following CLI commands wake up the AP from a deep-sleep mode:

```
(host) [mynode] #ap wake-up
ap-name <ap-name>
ip-addr <ip-addr>
ip6-addr <ip6-addr>
mac-list <mac_list>
wired-mac <wired_mac>
```

The following CLI commands display the APs in the deep-sleep mode. A new flag, p is introduced to show that the AP is in deep sleep mode:

```
(host) [mynode]# show ap database
AP Database
-----------
Name    Group    AP Type  IP Address       Status             Flags    Switch IP
Standby IP
----    -----    -------  ----------       ------             -----    ---------
----------
ap-205  default  205      191.191.191.252  Up 10d:8h:8m:6s     2p      192.192.189.1
0.0.0.0
ap-215  default  215      191.191.191.253  Up 33d:14h:1m:37s           192.192.189.1
0.0.0.0

Flags:  U = Unprovisioned; N = Duplicate name; G = No such group; L = Unlicensed
        I = Inactive; D = Dirty or no config; E = Regulatory Domain Mismatch
        X = Maintenance Mode; P = PPPoE AP; B = Built-in AP; s = LACP striping
        R = Remote AP; R- = Remote AP requires Auth; C = Cellular RAP;
        c = CERT-based RAP; 1 = 802.1X authenticated AP; 2 = Using IKE version 2
        u = Custom-Cert RAP; S = Standby-mode AP; J = USB cert at AP
        i = Indoor; o = Outdoor
        M = Mesh node; Y = Mesh Recovery
        z = Datazone AP
        p = In deep-sleep status

Total APs:2
```

The following CLI commands display all the pending APs in the per-md list, sends the AP_INFO AMON message for a particular AP, and tracks Green AP related counters:

```
(host) [mynode] show ap greenap
amon pending-ap {all | ip-addr <ip-addr> | ip6-addr <ip6-addr>}
counters{ap-name <ap-name> | ip-addr <ip-addr> | ip6-addr <ip6-addr> | wired-mac
<wired-mac>}
```

```
request pending-ap {all | ip-addr <ip-addr> | ip6-addr <ip6-addr>}
```



The amon command can only be run in the managed device.

The `show ap greenap counters` command can be used for debugging as it displays information such as such how many deep-sleep requests from Netinsight are received, and how many requests are dropped and so on.

Air Slice

Aruba's key RF differentiation, Air Slice, designed for 11ax APs optimizes user experience and assures QoS to enterprise applications. Air Slice combines AppRF and UCC for classifying applications and it also supports custom flow definitions. Air Slice uses a combination of priority queuing, dynamic WMM boosting, and 11ax based radio resource scheduling to prioritize enterprise applications in the presence of competing background traffic flows to meet latency and bandwidth requirements.



Air Slice is supported only on AP-534, AP-535 and AP-555 access points and it is mandatory to enable Deep Packet Inspection before configuring Air Slice.

The following procedure associates an Air Slice profile to an AP group profile:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles > AP > AP group**.
2. Select an AP group and select **Airslice**. To create a new profile, click **+.**
3. Enter a **Profile name**.
4. Select the **airslice-app monitoring** check-box to enable traffic monitoring for applications.
5. Select the **airslice-policy check-box** to optimize communication quality for applications.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following procedure describes how to enable reporting the clients application usage:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles > Controller Profile > Mgmt Config**.
2. Select a profile and enable the **AP application stats** checkbox.
3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

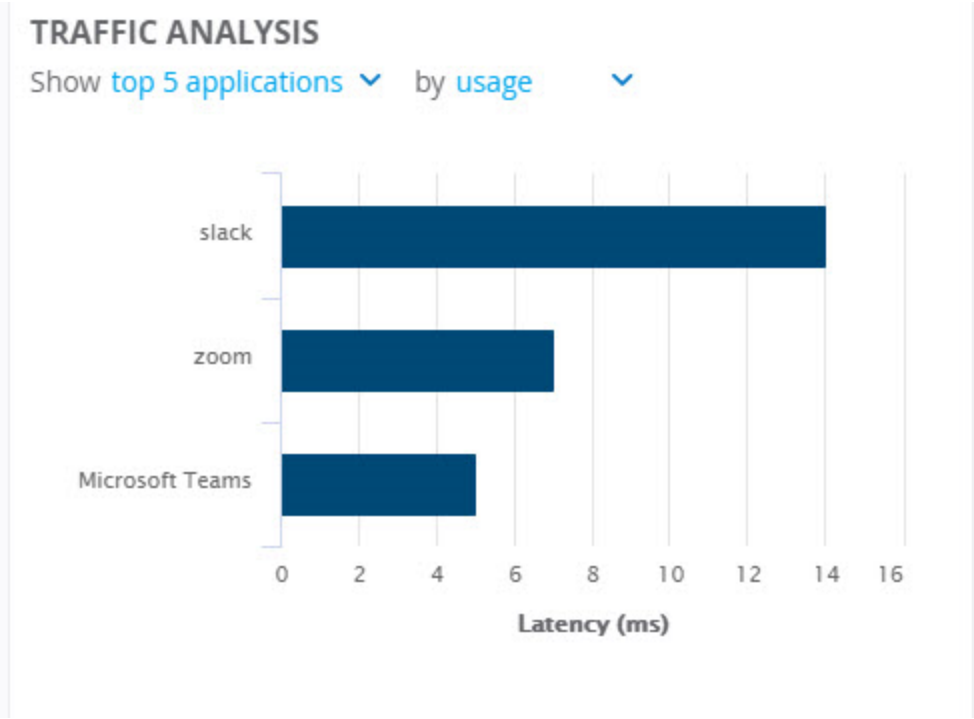
Monitoring Applications

The following procedure describes how to monitor applications used by the clients:

1. In the **Managed Network** node hierarchy, navigate to **Dashboard > Overview > Clients**.
2. Select a client. The **Traffic Analysis** pane displays the top 5 applications used by the client.
3. You can view the following information using the **by** drop-down list:
   - Usage
   - Latency

- Jitter
- Packet Loss

Figure 87 *Traffic Analysis Pane*



The following CLI commands configure an Air Slice profile:

```
(host) [mynode] (config) #airslice-profile <name>
(host) [mynode] (airslice profile <name>) airslice-app-monitoring
(host) [mynode] (airslice profile <name>) airslice-policy
(host) [mynode] (airslice profile <name>) clone
(host) [mynode] (airslice profile <name>) no
```

The following CLI commands enable Air Slice profile in **ap-group** and **ap-name** profiles:

```
(host) [mynode] ap-group <name>
(host) [mynode] (AP Group) #airslice-profile <name>
(host) [mynode] ap-name <name>
(host) [mynode] (AP name) #airslice-profile <name>
```

The following command enables reporting client application usage:

```
(host) [md] #mgmt-server profile <name>
(host) [md] (Mgmt Config profile) #ap-app-stats
```

Session ACLs for custom applications can be configured only using the CLI. The following command configures session ACL for a custom application:

```
(host) [mynode] (config) #ip access-list session airslice
(host) [mynode] (config-submode)#any host 1.2.3.4 tcp 34 permit markapp custom5
```

In addition to custom applications, Air Slice is supported only for the applications listed below:

- Zoom
- Slack

- Skype
- FaceTime
- WebEx
- GoToMeeting
- Office365
- Dropbox
- Amazon AWS
- Github
- Micrsoft Exel Online
- Onedrive
- Outlook
- Microsoft Planner
- Microsoft Powerpoint
- SharePoint
- Microsoft Sway
- Microsoft Teams
- Microsoft Word Online
- Yammer
- Wifi-Calling

Optional AP Configuration Settings

Once the AP has been installed and provisioned, you can use the WebUI or CLI to configure the optional AP settings described in the following sections:

- Spanning Tree on page 709
- PortFast on page 709
- AP Console Access Using a Backup ESSID on page 709
- Defining an RTLS Server on page 709
- AP Redundancy on page 710
- AP Maintenance Mode on page 710
- Energy Efficient Ethernet on page 710
- Smart Rate on page 711
- Configuring Energy Efficient Ethernet and Ethernet Link Speed on page 711
- Associating Ethernet Interface Link Profile with a Wired Port profile on page 712
- AP LEDs on page 713
- Suppressing Client Probe Requests on page 713
- BLE Operation Mode on page 714
- Intelligent Power Management on page 715
- Intelligent Thermal Management on page 717
- Intelligent Power and Temperature Monitoring on page 718
- Configuring the AP System Profile on page 718
- Configuring the AP Wired Port Profile on page 729
- Configuring the Dump Collection Profile

- Configuring Preferred Uplink on page 729
- Tri-Radio Mode for 550 Series Access Points on page 731

Spanning Tree

The STP can prevent loops in bridged Ethernet local area networks. Spanning tree settings can be configured via the WebUI and command-line interface.

To enable this feature, enable both the **Spanning Tree** parameter in the AP system profile and the **Spanning Tree**parameter in the AP wired port profile. For details, see Configuring the AP System Profile on page 718.

PortFast

The PortFast feature is introduced to avoid network connectivity issues. These issues are caused by delays in STP enabled ports moving from blocking-state to forwarding-state after transitioning from the listening and learning states. STP enabled ports that are connected to devices such as a single switch, workstation, or a server can access the network only after passing all these STP states. Some applications need to connect to the network immediately, else they will timeout.

Spanning Tree should be enabled on the access point before enabling PortFast. If PortFast is configured, it is enabled only on access mode ports and if PortFast-Trunk is configured, it is enabled on trunk-mode ports only. Only one of them can be set based on the port's switchport mode.

Enabling PortFast on an Access Port

Before enabling PortFast ensure that the switchport mode is set to **access**:

```
(host)[mynode] #show ap wired-port-profile <profile>
```

Execute the following commands in config mode to enable PortFast on an access port:

```
(host)[mynode] (config) #ap wired-port-profile "default"
(host)[mynode] (AP wired port profile "default") #portfast
```

Enabling PortFast on a Trunk Port

Before enabling PortFast ensure that the switchport mode is set to **trunk**:

```
(host)[mynode] #show ap wired-port-profile <profile>
```

Execute the following commands in config mode to enable PortFast on a trunk port:

```
(host)[mynode] (config) #ap wired-port-profile "default"
(host)[mynode] (AP wired port profile "default") #portfast-trunk
```

AP Console Access Using a Backup ESSID

This failover system allows users to access an AP console after the AP has disconnected from the managed device. By advertising backup ESSID in either static or dynamic mode, the user is still able to access and debug the AP remotely through a virtual AP. Settings for this feature are configured using the **Password for Backup**, **RF Band for Backup**, and **Operation for backup** parameters in the AP system profile. For details, see Configuring the AP System Profile on page 718.

Defining an RTLS Server

The RTLS server configuration enables the AP to send RFID tag information to an RTLS server. Currently, when configuring the RTLS server under **ap system-profile**, you can set the **station-message-frequency** parameter in the 1-3600 seconds range. Setting the frequency to 1 means a report is sent for

every station every second. A value of 5 means that a report for any particular station would be sent at 5 second intervals.

- Sending more frequent reports to the server can improve the accuracy of the location calculation.
- Configuring an AP to send reports more frequently adds additional load in terms of CPU usage.

Settings for this feature are configured using the **RTLS Server configuration** parameters in the **Advanced** section the AP system profile. For details, see Configuring the AP System Profile on page 718.

AP Redundancy

In conjunction with the managed device redundancy features described in Increasing Network Uptime With Redundancy Services on page 780 the information in this section describes redundancy for APs. Remote APs also offer redundancy solutions via a backup configuration, backup managed device list, and remote AP failback. For more information relevant to remote APs, see Remote Access Points on page 844.

The AP failback feature allows an AP associated with the backup managed device (backup LMS) to fail back to the primary managed device (primary LMS) if it becomes available.

If configured, the AP monitors the primary managed device by sending probes every 600 seconds by default. If the AP successfully contacts the primary managed device for the entire hold-down period, it will fail back to the primary managed device. If the AP is unsuccessful, the AP maintains its connection to the backup managed device, restarts the LMS hold-down timer, and continues monitoring the primary managed device.

Settings for this feature are configured using the LMS IP parameters in the **LMS settings** section of the AP system profile. For details, see Configuring the AP System Profile on page 718.

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 managed device 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.

The AP maintenance mode is configured by enabling **Maintenance Mode** parameter in the **Advanced** section of the AP system profile. For details, see Configuring the AP System Profile on page 718.

The following CLI commands display the status of APs in maintenance mode:

```
show ap config {ap-group <name>|ap-name <name>|essid <name>}
show ap debug system-status {ap-name <name>|bssid <name>| ip-addr <ipaddr>}
```

Energy Efficient Ethernet

Most new models of Aruba APs support the 802.3az Energy Efficient Ethernet standard, which allows the APs to consume less power during periods of low data activity. This setting can be enabled for provisioned APs or AP groups through the Ethernet Link profile. If this feature is enabled for an AP group, any APs in the group that do not support 802.3az will ignore this setting. For more information on configuring Energy Efficient Ethernet, see Configuring Energy Efficient Ethernet and Ethernet Link Speed on page 711.



802.3az is not supported on AP-215, AP-315, and 330 Series access points.

Thermal Shutdown Support in Access Points

Starting from ArubaOS 8.6.0.0, several Aruba AP models are enabled with thermal shutdown feature. The APs are equipped with one or more internal temperature sensors that initiate a shutdown when the AP's internal temperature crosses a set threshold. The AP then disconnects from the controller and operates in a low-power mode, allowing it to cool down. Once the AP reaches a normal operating temperature, it reconnects to the controller. This process of rebootstrap and reconnection is carried out for 5 times, until the connection is restored. If the connection between the AP and the controller still does not secure, the AP remains in the shutdown state till it is manually turned on.

Starting from ArubaOS 8.7.0.0, thermal shutdown for mesh mode APs is supported.

Smart Rate

HPE Smart Rate is a new multi-gigabit (1, 2.5, 5, 10 Gbps) twisted-pair network interface that is interoperable with the NBASE-T ecosystem of 2.5 or 5 Gbps products as well as with existing industry standard 1 GbE or 10 GbE devices. It allows the majority of existing cable installations found in campus LAN environments to provide higher bandwidth connectivity, distribute PoE power to connected devices, and secure the wired-link for next-generation 802.11ac applications.

With smart rate configuration enabled, an AP is capable of negotiating more than 1Gbps of link speed with a smart rate capable switch. 330 Seriesaccess points are capable of negotiating up to 5 Gbps speed. By default, the Ethernet interface speed is configured as **auto** (auto-negotiate) and the eth0 interface of 330 Series access points negotiate a 2.5 Gbps speed. To obtain 5 Gbps speed negotiation, enforce the speed value in the AP Ethernet Link profile. For more information on configuring the link speed, see Configuring Energy Efficient Ethernet and Ethernet Link Speed on page 711.

Configuring Energy Efficient Ethernet and Ethernet Link Speed

You can configure 802.3az or link speed using the WebUI or CLI.

The following procedure describes how to configure 802.3az EEE and Ethernet link speed:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** page.
2. Select **AP > AP Ethernet Link**, and select the Ethernet link profile you want to modify. The parameters for the profile are described in Ethernet Interface Link Profile Parameters on page 712.
3. The selected profile appears in the **Profile Details** window. The configuration parameters for the profile are described in Ethernet Interface Link Profile Parameters on page 712.
4. Select the **803.az** checkbox to enable energy efficient ethernet.
5. Select the required value from the **Speed** drop-down list to configure the Ethernet link speed. The speed value can be one of the following values:
   - **10**— 10 Mbps
   - **100**— 100 Mbps
   - **1000**—1 Gbps
   - **2500**— 2.5 Gbps
   - **5000**— 5 Gbps
   - **auto**— Auto-negotiate. This is the default value.
6. Save your changes.

The following CLI command enables support for 803.az EEE:

```
(host)[mynode] (config) #ap enet-link-profile <profile>
```

```
dot3az
```

The following CLI commands configure the Ethernet link speed:

```
(host) [mynode] (config) #ap enet-link-profile <profile>
(host) [mynode] (AP Ethernet Link profile "<profile>") #speed <speed>
```

where <speed> can take any of the following values:

- **10**— 10 Mbps
- **100**— 100 Mbps
- **1000**—1 Gbps
- **2500**— 2.5 Gbps
- **5000**— 5 Gbps
- **auto**— Auto-negotiate. This is the default value.

Table 130 describes the Ethernet Interface Link profile parameters.

Table 130: *Ethernet Interface Link Profile Parameters*

Parameter	Description
Speed	The speed of the Ethernet interface, either 10 Mbps, 100 Mbps, 1000 Mbps (1 Gbps), or auto-negotiated.
Duplex	The duplex mode of the Ethernet interface, either full, half, or auto-negotiated.
802.3az (EEE)	Select this check box to enable support for 802.1az Energy Efficient Ethernet.
Power Over Ethernet	Enable PoE for APs that support PoE.

Associating Ethernet Interface Link Profile with a Wired Port profile

By default, AP wired port profiles reference the Default Ethernet interface link profile. If you created a new Ethernet interface link profile to support 803.az, you can associate an AP wired port profile or Ethernet interface port configuration with the new Ethernet Interface link profile.

The following procedure describes how to associate a new Ethernet Interface link profile with an AP wired port profile.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** page.
2. Navigate to **AP > AP Wired Port Profile** on the **Profile** pane, then select the AP Wired Port profile you want to modify.
3. Click the Ethernet interface link profile currently associated with the AP wired port profile you want to modify. This profile appears below the AP Wired Port Profile in the **All Profiles** list.
4. Click the **Ethernet interface link profile** drop-down list at the top of the **Profile Details** window, and select a new Ethernet interface link profile.
5. Save your changes.

The following CLI commands associate a new Ethernet Interface link profile with an AP wired port profile.

```
(host)[node] (config) #ap wired-port-profile <profile>
```

```
enet-link-profile <profile>
```

AP LEDs

AP LEDs on 802.11n and 802.11ac APs 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. The AP LED mode is configured by enabling the **LED Operating Mode** parameter in the **General** section of the AP system profile. For details, see Configuring the AP System Profile on page 718.

Suppressing Client Probe Requests

The anyspot client probe suppression feature decreases network traffic by suppressing probe requests from clients attempting to locate and connect to other known networks. By reducing the frequency at which these messages are sent, this feature frees up network resources and improves network performance.

When an AP is configured to use this feature, the anyspot AP radio hides its configured ESSID in beacons, and compiles a list of other ESSIDs from detected neighboring APs. If the client sends a probe request without a specified ESSID, the anyspot AP will respond with a preconfigured ESSID.

When a client searches for a preferred network, that client sends the SSID of the preferred network in the probe request. The anyspot AP checks to see if there is a neighboring AP using that ESSID that can respond to the client request. If no matching network is found, the anyspot AP sends a response to the client using the SSID from the client request. If the client is authorized to connect to the anyspot AP, that client associates to AP. Once connected to the anyspot AP, the client recognizes the ESSID to which it is connected as one associated with its preferred network, and does not send out any further probe requests.



An AP radio can only use this feature when encryption is disabled. (That is, when the **operation mode** parameter in the AP radio WLAN SSSID profile is set to **opensystem**.)

You can define a list of excluded ESSIDs to which the anyspot AP will not respond. If a client sends probe request with an ESSID on the excluded ESSID list, the anyspot AP will not respond to the request, even if there is no neighboring AP using that ESSID. Excluded ESSIDs can be identified by exact name or a matching string.

The following procedure describes how to configure an anyspot profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** page.
2. Navigate to **Wireless LAN> Anyspot** on the **Profile** pane, then select the anyspot profile you want to modify.
3. Configure the anyspot parameters described in Table 131.

Table 131: *Anyspot Client Probe Suppression Configuration Parameters*

Parameter	Description
Enable Anyspot	Select this check box to enable the anyspot feature. Note that you must associate the anyspot profile with a virtual AP profile for the settings to take effect.

Table 131: *Anyspot Client Probe Suppression Configuration Parameters*

Parameter	Description
Exclude ESSID(s) (exact match)	An anyspot-enabled radio will not respond to client probe requests using an ESSID in the **Exclude ESSID** lists. To add an ESSID to the list, enter the full name of the ESSID, then click **Add**. To remove an ESSID from the list, select it and click **Delete**. ESSIDs from neighboring APs will automatically appear in this list as long as the anyspot-enabled AP can detect that ESSID.
Exclude ESSID(s) (containing string(s)	An anyspot-enabled radio will not respond to client probe requests using an ESSID in the **Exclude ESSID** list. To exclude ESSIDs that partially match a text string, enter that string then click **Add**. To remove a matching string from the list, select it and click **Delete**.
Preset ESSID(s)	The anyspot-enabled AP will not send an ESSID in beacons, but if a client sends a probe requests without ESSIDs (that is, the probe request is not looking for a specific network)then the anyspot-enabled AP will respond to the probe request with an ESSID from this list.

If you create a new Anyspot profile, use the procedure below to associate the anyspot profile with a selected WLAN via the virtual AP profile.

The following procedure describes how to associate a new Ethernet interface link profile with a wired port profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** page.
2. Select **AP > Virtual AP** on the **Profile** pane, then select the Virtual AP profile for the WLAN you want to modify.
3. Click the **Anyspot** profile currently associated with the Virtual AP profile. This profile appears below the Virtual AP Profile in the **All Profiles** list.
4. Click the **Anyspot profile** drop-down list and select the new **Anyspot** profile.

In the CLI

The following CLI commands to configure the anyspot profile, and associate an anyspot profile with a virtual AP:

```
(host)[node] (config) #wlan anyspot-profile <anyspot-profile>
(host)[node] (config) #wlan virtual-ap profile <profile>
anyspot <profile>
```

BLE Operation Mode

The **BLE Operation Mode** setting determines how the built-in BLE chip in the AP functions. You can configure this setting using the Mobility Master WebUI or CLI.



BLE functionality is enabled on 210 Series, 300 Series, 530 Series, and 550 Series FIPS APs.

This feature supports the following modes:

- **Beaconing**: The built-in BLE chip in the AP functions as an iBeacon combined with beacon management functionality.
- **Disabled**: The built-in BLE chip in the AP is turned off. This is the default setting.
- **DynamicConsole**: The built-in BLE chip in the AP functions as a regular iBeacon combined with beacon management functionality. However, when the link to the managed device is lost, the built-in

chip temporarily enables access to the AP console over BLE. This state of the BLE device may be rolled back to any of the other modes if the AP receives a different configuration setting for the **ble-op-mode** parameter from the new LMS.

- **PersistentConsole**: The built-in chip in the AP provides access to the AP console over BLE using a mobile application. This functionality is the superset of the **Beaconing** mode.

Settings for this feature are configured using the **BLE Operation Mode** parameter in the **Advanced** section the AP system profile. For details, see Configuring the AP System Profile on page 718.

Intelligent Power Management

The Intelligent Power Management (IPM) feature actively measures the power utilization of the AP and dynamically adapts to the power resources. IPM dynamically limits the power requirement of the AP as per the available power resources. This is in contrast to the existing static power management method where the power profiles such as POE-AF, POE-AT, POE-DC, or LLDP are hard-coded for each AP. The reduction steps and the associated priority values are configured to control the AP power consumption within the power budget.



IPM is supported on all AP platforms except 203H Series, 203R Series, 207 Series, 303 Series, 303P Series, 318 Series, 320 Series, 360 Series, and 370 Series access points.

The following procedure configures IPM:

1. In the **Managed Network** node hierarchy, navigate to **Configuration** > **System** > **Profiles** tab.
2. In the **All Profiles** list, expand the **AP** menu, then select **AP system**.
3. Select the AP system profile that you want to edit, or click **+** to create a new profile.
4. Expand **IPM Configuration**, and configure the following parameters:
   - **IPM activation**—Select the check box to enable IPM.
   - **IPM power reduction steps with priorities**—Click **+** to set the IPM power reduction steps and specify their priorities. Configure the following parameters in the **Add New** pop-up window:
     - **IPM_step_priority**—Enter a value to define the priority of the IPM power reduction step.

       A lower value implies the highest priority, and is implemented first over a priority with higher value.

     - **IPM_step**—Select the desired IPM reduction step as described inTable 132.
5. Click **OK**.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI command enables IPM.

```
(host)[mynode] (config)# ap system-profile default
(host)[mynode](AP system profile "default")# ipm-enable
```

Configuring Reduction Steps

The reduction steps and their order are configured using either the **IPM power reduction steps with priorities** parameter in the WebUI, or the **ipm-power-reduction-step-prio ipm-step <reduction**

step> priority <priority> parameter in the **ap system-profile** command. The reduction steps are associated with priority values, and the priority settings are defined by integer values, where the lower values have the highest priority and are implemented first. When triggered, the reduction steps are applied sequentially starting with the reduction step that is assigned the highest priority value. The reduction steps are applied only when the AP exceeds the power budget or threshold temperature.

The following table describes the reduction steps and their corresponding WebUI and CLI parameters.

Table 132: *IPM Reduction Step Parameters*

Reduction Step	WebUI/CLI Parameter
Reduce the CPU frequency to 25%.	`cpu_throttle_25`
Reduce the CPU frequency to 50%.	`cpu_throttle_50`
Reduce the CPU frequency to 75%.	`cpu_throttle_75`
Disable the second Ethernet port.	`disable_alt_eth`
Disable the PSE port.	`disable_pse`
Disable the USB port.	`disable_usb`
Reduce the 2.4 GHz radio chains to 1x1.	`radio_2ghz_chain_1x1`
Reduce the 2.4 GHz radio chains to 2x2.	`radio_2ghz_chain_2x2`
Reduce the 2.4 GHz radio chains to 3x3.	`radio_2ghz_chain_3x3`
Reduce the power of 2.4 GHz radio by 3 dB from the maximum value.	`radio_2ghz_power_3dB`
Reduce the power of 2.4 GHz radio by 6 dB from the maximum value.	`radio_2ghz_power_6dB`
Reduce the 5 GHz radio chains to 1x1.	`radio_5ghz_chain_1x1`
Reduce the 5 GHz radio chains to 2x2.	`radio_5ghz_chain_2x2`
Reduce the 5 GHz radio chains to 3x3.	`radio_5ghz_chain_3x3`
Reduce the 5 GHz radio chains to 4x4.	`radio_5ghz_chain_4x4`
Reduce the 5 GHz radio chains to 5x5.	`radio_5ghz_chain_5x5`
Reduce the 5 GHz radio chains to 6x6.	`radio_5ghz_chain_6x6`
Reduce the 5 GHz radio chains to 7x7.	`radio_5ghz_chain_7x7`
Reduce the power of 5 GHz radio by 3 dB from the maximum value.	`radio_5ghz_power_3dB`
Reduce the power of 5 GHz radio by 6 dB from the maximum value.	`radio_5ghz_power_6dB`



To reduce the CPU power gradually, the smallest reduction is allocated a higher priority value so that the minimum reduction step is implemented first. For example, the **cpu_throttle_50** parameter should have a higher priority value than the **cpu_throttle_25** parameter, so that IPM gradually reduces the CPU throttle or power usage based on the priority list.

The following CLI commands configure the IPM reduction steps and their priority.

```
(host)[mynode] (config)# ap system-profile default
(host)[mynode](AP system profile "default")# ipm-power-reduction-step-prio ipm-step
disable_usb priority 1
(host)[mynode](AP system profile "default")# ipm-power-reduction-step-prio ipm-step
radio_2ghz_chain_3x3 priority 2
(host)[mynode](AP system profile "default")# ipm-power-reduction-step-prio ipm-step
radio_5ghz_power_3dB priority 3
(host)[mynode](AP system profile "default")# ipm-power-reduction-step-prio ipm-step
cpu_throttle_75 priority 4
```

The following CLI command verifies the IPM configuration.

```
(host)[mynode] (config)# show running-config
...
ipm
ipm-power-reduction-step-prio ipm-step disable_usb priority 1
ipm-power-reduction-step-prio ipm-step radio_2ghz_chain_3x3 priority 2
ipm-power-reduction-step-prio ipm-step radio_5ghz_power_3dB priority 3
ipm-power-reduction-step-prio ipm-step cpu_throttle_75 priority 4
```

Intelligent Thermal Management

The Intelligent Thermal Management (ITM) feature measures the internal temperature of the AP and dynamically adapts operations to reduce the internal temperature. When the internal temperature of the AP exceeds the maximum threshold temperature, the operations of the AP are throttled down to reduce its internal temperature.

The reduction steps applied for ITM are configured using either the **IPM power reduction steps with priorities** parameter in the WebUI, or the **ipm-power-reduction-step-prio ipm-step <reduction step> priority <priority>** parameter in the **ap system-profile** command.



ITM is supported only on AP-505H, AP-518, and 570 Series access points.

The following procedure configures ITM:

1. In the **Managed Network** node hierarchy, navigate to **Configuration** > **System** > **Profiles** tab.
2. In the **All Profiles** list, expand the **AP** menu, then select **AP system**.
3. Select the AP system profile that you want to edit, or click **+** to create a new profile.
4. Expand **ITM Configuration**, and select the **ITM activation** check box.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI command enables ITM.

```
(host)[mynode] (config)# ap system-profile default
(host)[mynode](AP system profile "default")# itm-enable
```

Intelligent Power and Temperature Monitoring

Intelligent Power and Temperature Monitoring (IPTM) is a combination of Intelligent Power Management (IPM) and Intelligent Thermal Management (ITM) features that optimize AP operations in changing power and temperature conditions. IPTM adapts the performance of the AP to provide consistent network experience according to the power usage and operating temperature of the AP.

To manage this optimization, you must configure a set of reduction steps and associate them with a priority value. IPTM applies the sequence of reduction steps as defined by the associated priority value until the AP starts functioning within the power budget and threshold temperature. This happens dynamically as IPTM constantly monitors the power consumption and temperature of the AP and applies the corresponding reduction steps if the AP exceeds the power and temperature threshold.

For more information on IPM and ITM, see the following topics:

- Intelligent Power Management on page 715
- Intelligent Thermal Management on page 717

Points to Note

- IPM and ITM must be configured separately using the WebUI or the CLI.
- IPM and ITM are disabled by default.
- IPM must be enabled for ITM to function.
- The reduction steps applied for IPM and ITM are the same and are configured using either the **IPM power reduction steps with priorities** parameter in the WebUI, or the **ipm-power-reduction-step-prio ipm-step <reduction step> priority <priority>** parameter in the **ap system-profile** command.
- Initially, all functionalities of the AP are turned on by IPM and ITM. The reduction steps are triggered only when the power consumption of the AP exceeds the threshold associated with the power budget or when the temperature of the AP goes beyond the threshold value.
- IPM and ITM do not override pre-existing settings that restrict AP functionality. For example, when the USB port is disabled in the provisioning profile, the AP does not enable the functionality when the reduction steps are revoked.

Configuring the AP System Profile

The AP system profile configuration settings are divided into four groups, **General**, **LMS Settings**, **Remote AP** and **Advanced**. The **General**, **LMS Settings**, and **Remote AP** sections of this profile include configuration settings that often need to be adjusted to suit a specific network. The **Advanced** tab includes settings that do not need frequent adjustment or should be kept at their default values.

The AP system profile can be enabled using SSH for high-end encryption. The AP provides a remote login function for each user to access the AP shell, SSH in such cases enables state-of-art encryption to avoid network attack or risk of malicious users. When an Aruba AP user establishes a remote login function using Linux, the AP uses Telnet to establish the connection. Now an Aruba AP can be connected using SSH protocol for security and high-end encryption.



For console access via SSH, the user name is **root** and the password will the be same the console password specified in the AP system profile.

The following procedure describes how to configure AP settings using the AP system profile:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles** tab.
2. In the **All Profiles** list, expand the AP menu, then select **AP system**.
3. Select the AP system profile you want to edit, or click + to create a new profile.
4. Configure the profile parameters described in AP System Profile Configuration on page 719, then click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy Changes**.



To specify the server details to receive the core dump when an AP process crashes, select an AP system profile and click on **Dump Collection**. To allow the core dump files to be sent to the managed device, access the managed device command-line interface and issue the **ap-crash-transfer** command.

Table 133: *AP System Profile Configuration*

Parameter	Description
Profile Name	
General	
RF Band	For APs that support both 802.11a and 802.11b/g RF bands, specify the RF band in which the AP should operate: ■ g = 2.4 GHz ■ a = 5 GHz
RF Band for AM Mode scanning	For Air Monitors that support both 802.11a and 802.11b/g RF bands, specify the RF band which the AM should scan: ■ a = 5 GHz ■ all = both radio bands ■ g = 2.4 GHz
Native VLAN ID	Native VLAN for bridge mode virtual APs (frames on the native VLAN are not tagged with 802.1q tags).
WIDS AMPDU Optimization	Reduce the number of frames copied for the purpose of WIDS aggregate MPDU Optimization. Default: Enabled.
Session ACL	Session ACL configured with the **ip access-list session** command. **NOTE:** This parameter requires the PEFNG license.
Corporate DNS Domain	Name of domain that is resolved by corporate DNS servers. Use this parameter when configuring split-tunnel forwarding.
SNMP sysContact	SNMP system contact information.
LED operating mode	The operating mode for the LEDs on 802.11n-capable indoor AP. supported options are normal mode, and off, which disables all LEDs.

Table 133: *AP System Profile Configuration*

Parameter	Description
LED override	Override the LED action for single-LED APs in normal LED operating mode. If enabled, this feature disables the LED auto-turn-off function.
Driver log level	Level of AP driver logs sent to the syslog server. Supported options are: ▪ **emergencies**: System is unusable ▪ **alerts**: Immediate action needed ▪ **critical**: Critical Conditions ▪ **errors**: Error Conditions ▪ **warnings**: Warning conditions ▪ **notifications**: Normal but significant conditions ▪ **informational**: Informational Messages ▪ **debugging**: Debugging Messages
Console log level	Level of AP console logs sent to the AP console. Supported options are: ▪ **emergencies**: System is unusable ▪ **alerts**: Immediate action needed ▪ **critical**: Critical Conditions ▪ **errors**: Error Conditions ▪ **warnings**: Warning conditions ▪ **notifications**: Normal but significant conditions ▪ **informational**: Informational Messages ▪ **debugging**: Debugging Messages **NOTE:** The default option of this feature is **emergencies**. Do not change the console log level without prior supervision from the Aruba Technical Support team.
SAP MTU	Maximum Transmission Unit, in bytes, on the wired link for the AP.
RAP MTU	Configures the maximum size of the GRE packets exchanged between a RAP and the managed device.
Flex Radio Mode	Set the mode for flexible radios in 2.4 GHz, 5 GHz, and dual mode.
Dual 5GHz Mode	Set the mode for dual 5 GHz APs. This parameter is disabled by default.
Secondary Master IP/FQDN	The secondary Mobility Master is configured to be used when a RAP is not able to reach the primary Mobility Master.
Spanning Tree	Enable the spanning-tree protocol.
AP multicast aggregation	Enable multicast aggregation at AP

Table 133: *AP System Profile Configuration*

Parameter	Description
AP ARP attack protection	Drop ARP packets coming from wired or wireless clients with AP gateway IP address. In other words, disallow ARP attack from un-trusted ports.
AP multicast aggregation allowed VLANs	Enable a list of VLANs where AP multicast aggregation is allowed.
Advanced	***Advanced***
Tunnel heartbeat interval	Set the interval between heartbeat messages between a remote or campus AP and its associated managed device. An increase in the heartbeat interval increases the time it will take for an AP to detect the loss in connectivity to the managed device, but can reduce internet bandwidth consumed by a remote AP. The supported range is 1-60 seconds, and the default value is 1 second.
LMS ping interval	Specifies the interval at which application level ping needs to be sent to primary managed device to check the reachability. Applicable only for RAP. **NOTE:** If this parameter is changed, UDP session timeout on an intermediate router which performs NATing should be set accordingly. The preferred timeout value is (lms-ping-interval + 30sec). The supported range is 10-60 seconds, and the default value is 20 seconds.
Bootstrap threshold	Number of consecutive missed heartbeats on a GRE tunnel (heartbeats are sent once per second on each tunnel) before an AP rebootstraps. On the managed device, the GRE tunnel timeout is 1.5 x bootstrap-threshold; the tunnel is torn down after this number of seconds of inactivity on the tunnel. The supported range is 1-65535, and the default value is 8.
Double Encrypt	This parameter applies only to remote APs. Use double encryption for traffic to and from a wireless client that is connected to a tunneled SSID. When enabled, all traffic is re-encrypted in the IPsec tunnel. When disabled, the wireless frame is only encapsulated inside the IPsec tunnel. All other types of data traffic between the managed device and the AP (wired traffic and traffic from a split-tunneled SSID) are always encrypted in the IPsec tunnel.
Heartbeat DSCP	Assign a DSCP value to AP heartbeats to prioritize heartbeats traveling over low-speed links. The supported range is 0-63, and the default value is 0. For more information, see Validating and Optimizing AP Connectivity on page 645.
Maintenance Mode	Enable or disable AP maintenance mode. This setting is 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 when deploying, maintaining, or upgrading the network. The managed device still generates debug syslog messages if debug logging is enabled.

Table 133: *AP System Profile Configuration*

Parameter	Description
Maximum Request Retries	Maximum number of times to retry AP-generated requests, including keepalive messages. After the maximum number of retries, the AP either tries the IP address specified by the **bkup-lms-ip** (if configured) or reboots.
Recovery Mode	Select either the legacy recovery mode or the auto mode (fast recovery). ▪ **Legacy**— On detecting a firmware assert, the AP transfers the coredump to the managed device and executes an AP reboot. ▪ **Auto**—On detecting a firmware assert, the AP executes the fast recovery process in the radio affected instead of rebooting the AP. This reduces the downtime of the AP in the network. If the AP detects a core dump with a valuable information during a firmware assert, then it transfers the core dump to the managed device and the AP reboots. This is the default mode.
Request Retry Interval	Interval, in seconds, between the first and second retries of AP-generated requests. If the configured interval is less than 30 seconds, the interval for subsequent retries is increased up to 30 seconds.
Number of IPSEC retries	Number of times the AP will try to create an IPsec tunnel with the Mobility Controller before the AP will reboot. If you specify a value of 0, and AP will not reboot if it cannot create the IPsec tunnel. The supported range of values is 0-1000 retries, and the default value is 85 retries.
Root AP	Defines a remote AP as the root AP in a branch network with a multi-AP hierarchy.
AeroScout RTLS Server	Enables the AP to send AeroScout tag information to an RTLS server. You must specify the IPv4/IPv6 address or DNS server and port number of the server to which location reports are sent. RTLS station reporting includes information for APs and the clients that the AP has detected. If you select the **Include Unassociated Stations** option, the station reports will also include information about clients not associated to any AP. By default, unassociated clients are not included in station reports.
IP or DNS	IPv4/IPv6 address or the DNS of the AeroScout server to which location reports are sent.
Port	Port number on the AeroScout server to which location reports are sent.
includeUnassocSta	If you select the **Include Unassociated Stations** option, the station reports will also include information about clients not associated to any AP. By default, unassociated clients are not included in station reports.
RTLS Server configuration	Enables the AP to send RFID tag information to an RTLS server. You must specify the IP address or DNS server and port number of the server to which location reports are sent, a shared secret key, and the frequency at which packets are sent to the server.

Table 133: *AP System Profile Configuration*

Parameter	Description
	RTLS station reporting includes information for APs and the clients that the AP has detected. For more information on configuring RTLS server configuration, see Defining an RTLS Server on page 709.
SES-imagotag ESL Server IP	SES-imagotag ESL Server IP configures the IP Address of ESL Server. Adding server IP addresses allows managing and controlling bulk servers at the same time.
SES-imagotag ESL Channel	Configures the channel of SES-imagotag ESL Radio. You can select a radio channel ranging from 0 to 10. These channels do not correspond to standard 802.11 channels. **NOTE:** There are 11 pre-defined, independent radio channels that you can configure. The recommended channels are 3, 5, 8, 9, and 10 as they connect faster.
IP or DNS	IPv4/IPv6 address or the DNS of the RTLS server to which location reports are sent.
Port	RTLS server port number
frequency	Specify how often to send station RSSI update messages to the server. The supported range is 1-3600 seconds, and the default setting is 30 seconds.
key	Shared secret key for the RTLS server.
Include Unassociated Stations	If you select the **Include Unassociated Stations** option for an RTLS server, the station reports will also include information about clients not associated to any AP. By default, unassociated clients are not included in station reports.
RTLS Server Compatibility Mode	The compatibility mode controls the format of tag frames forwarded to the RTLS server. Enabling this mode will enable legacy format (includes a 2 byte padding), and disabling this mode will remove the padding. The tag frame format will be the same across all AP models. This feature is enabled by default
Slow Timer Recovery by rebooting itself	If you enable this option, ArubaOS checks for a slow CPU timer, and if it detects an issue, restarts the AP without logging a reason for the reboot. This feature is supported on RAP-108/ RAP-109 access points.
Telnet/SSH	Select this check box to enable telnet or SSH to the AP. **NOTE:** Telnet is enabled on an AP running ArubaOS 8.6.0.0 or previous versions. SSH is enabled on an AP running ArubaOS 8.7.0.0.
Console Enable	Enable console port on the AP.
AP Console Protection	Enable the AP console protection by requiring a password to get AP console access.

Table 133: *AP System Profile Configuration*

Parameter	Description
AP Console Password	Sets the AP console password on the controller. If configured, you must enter this password to get AP console access. If not configured, the controller generates a default random password which can be viewed by executing the **encrypt disable** command followed by the **show ap system-profile <profile-name>** command. To disable the AP console password, execute the **shell-passwd passworddisabled** command. **NOTE:** The **passworddisabled** is a special character string. On entering this string, the controller disables the AP console password.
Password for Backup	Allows client access to adjust the band and mode settings for the backup ESSID.
AP USB Power mode	Enables or disables the USB port on AP platforms that have external USB ports. The supported values are as follows: ▪ **auto**: Detects USB power mode automatically ▪ **disable**: Disables USB power ▪ **enable**: Enables USB power **NOTE:** This parameter is applicable to AP-214, AP-215, AP-224, AP-225, AP-205H, AP-303H, AP-304, AP-305, AP-314, AP-315, AP-324, AP-325, AP-334, AP-335, AP-344, AP-345, AP-203R, and AP-203RP access points only.
AP POE mode	Choose POE mode on the AP platforms with dual POE support.
RF Band for Backup	Band on which the controller broadcasts the backup ESSID. Supported values are as follows: ▪ **a**: 802.11a ▪ **all**: all bands. This is the default setting. ▪ **g**: 802.11g
Operation for Backup	This parameter allows AP console access using a backup ESSID, allowing users to access an AP console after the AP has disconnected from the controller. When the AP advertises a backup ESSID in either static or dynamic mode, a user is able to access and debug the AP remotely through a virtual AP. The default setting for this feature is **off**. Select **dynamic** or **static** to enable this feature and select the mode by which the controller broadcasts the backup ESSID.
BLE Endpoint URL	URL of the Meridian server to which the BLE sends monitoring data.
BLE Auth Token	The BLE endpoint authorization token is a text string of 1-255 characters used by the BLE to authorize to and securely communicate with the Beacon Management Console. This token is unique for each deployment.
BLE Operation Mode	Determines how the built-in BLE chip in the AP functions. BLE chip can be in one of the following four modes:

Table 133: *AP System Profile Configuration*

Parameter	Description
	▪ **Beaconing**: The built-in BLE chip in the AP functions as an iBeacon combined with beacon management functionality. ▪ **Disabled**: The built-in BLE chip in the AP is turned off. This is the default setting. ▪ **Dynamic Console**: The built-in BLE chip in the AP functions as a regular iBeacon combined with beacon management functionality. However, when the link to the controller is lost, the built-in chip temporarily enables access to the AP console over BLE. This state of the BLE device may be rolled back to any of the other modes if the AP receives a different configuration setting for the **ble-op-mode** parameter from the new LMS. ▪ **Persistent Console**: The built-in BLE chip in the AP provides access to the AP console over BLE using a mobile application. This functionality is the superset of the **Beaconing** mode. **NOTE:** BLE is disabled for ArubaOS FIPS.
GRE Offload	
Health Check	The AP Health check feature uses ping probes to check reachablility and latency levels for the connection between the AP and the controller.
mode	Ping probe mode is the only mode currently supported by this feature.
packetsize	The size, in bytes, of a ping datagram. The supported range of values is 10-2000.
burstcnt	Number of probes to be sent during the probe frequency interval defined by the **frequency** health-check parameter. The supported range of values is 1-16.
freq	Probe interval, in seconds. The WAN health-check feature sends the number of probes defined by the **burst-size** parameter during each frequency interval defined by this **frequency** parameter. The supported range of values is 10-300.
report	Number of seconds between health check reports sent from the AP to the controller. The supported range of values is 60-3600.
retrycnt	Number of times the attempts to resend a probe. The supported range of values is 1-10 retry attempts.
AirMatch Report Period	Change the frequency period which AirMatch starts measuring the RF environment. The default value is 30 minutes, and the supported range of values is 5-180 minutes.

Table 133: *AP System Profile Configuration*

Parameter	Description
AirMatch Measurement Duration	Change the AirMatch RF measurement duration from the default value of five minutes to any value from 5-60 minutes. A value of 0 disables AirMatch RF environment measurements.
AirMatch Report Enabled	Enable or disable AirMatch reports. Each AP in a Mobility Master deployment measures its RF environment for a five minute duration, every 30 minutes. Mobility Master uses this information to compute an optimal solution, then deploys the latest RF plan by sending updated settings to the APs every 24 hours. This feature is enabled by default.
LMS	
LMS IP	This parameter specifies the IP address of the LMS—the managed device—which is responsible for terminating user traffic from the APs, and processing and forwarding the traffic to the wired network. This can be the IP address of the managed device or Mobility Master. When using redundant managed devices as the LMS, set this parameter to be the VRRP IP address to ensure that APs always have an active IP address with which to terminate sessions. **NOTE:** If the LMS-IP is blank, the access point will remain on the managed device that it finds using methods like DNS or DHCP. If an IP address is configured for the LMS IP parameter, the AP will be immediately redirected to the managed device at that address.
Backup LMS IP	This parameter specifies the IP address of a *backup* to the IP address specified with the **lms-ip** parameter.
LMS IPv6	This parameter specifies the IPv6 address of the LMS —the managed device—which is responsible for terminating user traffic from the APs, and processing and forwarding the traffic to the wired network. This can be the IP address of the managed device or Mobility Master. When using redundant managed devices as the LMS, set this parameter to be the VRRP IP address to ensure that APs always have an active IP address with which to terminate sessions.
Backup LMS IPv6	This parameter specifies the IPv6 address of a *backup* to the IPv6 address specified with the lms-ipv6 parameter.
LMS Preemption	When this parameter is enabled, the AP automatically reverts to the primary LMS IP address when it becomes available.
LMS Hold-down Period	Time, in seconds, that the primary LMS must be available before an AP returns to that LMS after failover.
Remote AP	
Remote-AP DHCP Server VLAN	VLAN ID of the remote AP DHCP server used if the managed device is unavailable. This VLAN enables the DHCP server on the AP (also known as the remote AP DHCP server VLAN). If you enter the native VLAN ID, the DHCP server is unavailable.
Remote-AP DHCP Server ID	IP address used as the DHCP server identifier.

Table 133: *AP System Profile Configuration*

Parameter	Description
Remote-AP DHCP Default Router	IP address for the default DHCP router.
Remote-AP DHCP DNS Server	IP address of the DNS server.
Remote-AP DHCP Pool Start	Configures a DHCP pool for remote APs. This is the first IP address of the DHCP pool.
Remote-AP DHCP Pool End	Configures a DHCP pool for remote APs. This is the last IP address of the DHCP pool.
Remote-AP DHCP Pool Netmask	Configures a DHCP pool for remote APs. This is the netmask used for the DHCP pool.
Remote-AP DHCP Lease Time	The amount of days that the assigned IP address is valid for the client. Specify the lease in <days>. A value of 0 indicates the IP address is always valid; the lease does not expire.
Remote-AP uplink total bandwidth	This is the total reserved uplink bandwidth (in Kilobits per second).
Remote-AP bw reservation 1 **Remote-AP bw reservation 2** **Remote-AP bw reservation 3**	Session ACLs with uplink bandwidth reservation in kilobits per second. You can specify up to three session ACLs to reserve uplink bandwidth. The sum of the three uplink bandwidths should not exceed the Remote-AP uplink total bandwidth.
Remote-AP Local Network Access	Enable or disable local network access across VLANs in a Remote-AP.
IPM Configuration	
IPM Activation	Enables the IPM system for power management.
IPM power reduction steps with priorities	Sets up all the IPM power reduction steps. Select the **Ipm_step_priority** and **IPM_step** to set the power reduction priority.
ITM Configuration	
ITM Activation	Enables the intelligent temperature monitoring system. If enabled, the temperature is dynamically controlled and the AP is allowed to cool down. **NOTE:** This feature can be enabled only when the IPM feature is enabled. This feature is supported only on 570 Series access points.

The following command configures the AP system profile in the command-line interface:

```
(host)[mynode](config)#ap system-profile <profile>
```

The following command configures the recovery mode parameter:

```
(host) [mynode]  (config) #ap system-profile <profile-name>
(host) [mynode] (AP system profile "<profile-name>") #recovery-mode legacy
```

The following **show** command displays the new statistics on AP recovery mode:

```
(host)[mynode] #show ap details ap-name <ap-name>
AP "<ap-name>" Radio 0 Operating Information
```

```
----------------------------------------
Item                  Value    Source
----                  -----    ------
Very High throughput  Enabled  Configuration
High throughput       Enabled  Configuration
Mode                  AP       Configuration
Band                  802.11a
Primary Channel       36       Configuration
80MHz Channel Group   36-48    Configuration
EIRP                  10.0     Configuration
Fast recovery start  2017-03-09 11:57:56
Fast recovery end    2017-03-09 11:58:00
Fast recovery         1
```

AP Hardware Offload

Hardware offload is a technology used by supporting AP models to achieve high throughput without occupying CPU resource. This feature is enabled by default in the AP system profile. You can disable it in the AP system profile for test or debugging purpose only, under the direct supervision of Aruba Support.

The following table describes the AP hardware offload capabilities.

Table 134: *AP Hardware Offload Capability*

Hardware Offload Item	Description	Supported APs
GRE Offload	Offloads GRE tunnel into hardware to achieve high throughput without utilizing CPU resource. This parameter is used for Campus APs or Remote APs in tunnel mode or D-tunnel mode of virtual AP traffic management profile.	320 Series, 330 Series, 530 Series, and 550 Series access points.
Bridge Offload	Offloads TCP client or UDP traffic into hardware. This parameter is used for bridge mode virtual AP or wired APs.	530 Series and 550 Series access points.

The following procedure configures AP hardware offload using the AP system profile.

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles** tab.
2. In the **All Profiles** list, expand the AP menu, then select **AP system**.
3. Select the AP system profile you want to edit, or click + to create a new profile.
4. Under **Advanced**, select the **GRE Offload** or **Bridge Offload** check box, and click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following commands configure hardware offload parameters in the AP system profile.

```
(host)[mynode](config)# ap system-profile sys-635
(host)[mynode](AP system profile "sys-635")# gre-offload
(host)[mynode](AP system profile "sys-635")# bridge-offload
```

AP Fast Recovery

The 300 Series, 530 Series, and 550 Series APs provide support for the AP Fast Recovery feature. On detecting a firmware assert, the AP executes the fast recovery process in the affected radio. This avoids

rebooting of the AP unnecessarily, thereby reducing the downtime of the AP in the network. If the AP detects a core dump with a valuable information during a firmware assert, then it transfers the core dump to the managed device and reboots. See AP System Profile Configuration on page 719 for more details.

SES-imagotag ESL System

Starting from ArubaOS 8.4.0.0, AP-303H, 300 Series access points, 310 Series access points, 320 Series access points, 330 Series access points, 340 Series access points, and 510 Series access points provide support for SES-imagotag's Electronic Shelf Label (ESL) system. ESL is used by various retailers to display the price and other associated information for products kept on retail shelves. SES-imagotag's ESL system enable APs to configure ESL-Radio, label, ESL-Server, and Client Software. The ESL-Radio is a USB dongle that works on the 2.4G frequency band. The ESL-Server is a management system that controls product labeling and the Client Software is the control center for all ESL-Servers. These centers help in controlling and executing various tasks such as changing images to labels, assigning tags, resetting labels, refreshing displays, switching to preloaded pages, etc. By enabling and using an ESL system, retail labeling becomes easier and efficient. Aruba APs integrated with SES-imagotag enables access to WIFI and ESL services simultaneously. You can set a server IP and channel range from the advanced settings in AP profiles. See SES-imagotag ESL System on page 729 for more details.

Configuring Preferred Uplink

Starting ArubaOS 8.4.0.0 ethernet port1 can be configured as the primary uplink and ethernet port0 can be configured as the downlink interface, in an active-standby uplink mode of deployment. This enhancement is supported in AP-318, AP-374, AP-375, AP-377.

The following CLI commands configure ethernet port1 as the primary uplink:

```
(host) [mynode] (config) #provision-ap
(host) [mynode] (config-submode)#read-bootinfo
(host) [mynode] (config-submode)#read-bootinfo ap-name ap_318
(host) [mynode] (config-submode)#preferred_uplink
(host) [mynode] (config-submode)#preferred_uplink eth1
```

Configuring the AP Wired Port Profile

This profile is only applicable to APs with Ethernet ports. Use this profile to enable or disable the wired port, define an AAA profile for wired port devices, and associate the port with an Ethernet link profile that defines its speed and duplex values.

The following procedure describes how to configure the AP wired port profile:

1. Navigate to the **Configuration > System > Profiles** page.
2. Select **AP > AP Wired Port**, and select the AP wired port profile you want to modify. The parameters for the profile are described in Ethernet Interface Link Profile Parameters on page 712.

Table 135: *AP Wired Port Profile Parameters*

Parameter	Description
Shut down	Disable the wired AP port.

Parameter	Description
Remote AP Backup	Enable this option to use the wired port on a Remote AP for local connectivity and troubleshooting when the AP cannot reach the managed device. If the AP is not connected to the managed device, no firewall policies will be applied when this option is enabled. (The AAA profile will be applied when the AP is connected to managed device).
Bridge Role	Role that is assigned to a user if split-tunnel authentication fails.
Time to wait for authentication to succeed	Authentication timeout value, in seconds, for devices connecting the AP's wired port. The supported range is 1-65535 seconds, and the default value is 20 seconds.
Spanning Tree	Enables the spanning-tree protocol.
Portfast	Enables portfast for AP wired access ports. Spanning tree must be enabled before this command can be used.
Portfast on trunk	Enables portfast for AP wired trunk ports. Spanning tree must be enabled before this command can be used.
Loop Protect Enable	Enables loop protection on AP wired ports.
Loop Detection Interval	Time, in seconds, to send loop detection packet. The supported range is 1 to 10 seconds and the default value is 2 seconds.
Storm Control Broadcast	Enables storm control broadcast. If the number of broadcast packets per second on one port in the AP exceeds the configured threshold, the port is shutdown.
Storm Control Broadcast Threshold:	Storm control broadcast threshold in packets per second after which the port is shutdown. The default value is 2000 packets per second.
Auto Recovery Enable	Enables automatic recovery of the port in the AP that is shut down because of loop protection. After the automatic recovery, if the loop re-occurs, then the port is shut down again.
Auto Recovery Interval	Time, in seconds, to automatically recover the port in the AP that is shut down because of loop protection. The supported range is 30 to 43200 seconds and the default value is 300 seconds.

The following CLI command configures the AP wired port profile:

```
(host)[node] (config) #ap wired-port-profile <profile>
```

Configuring the Dump Collection Profile

The dump collection profile configures the settings for collecting the core dump when an AP process crashes.

The following procedure describes how to configure dump collection profile settings:

1. In the **Managed Network** node hierarchy, navigate to **Configuration** > **System** > **Profiles** tab.
2. In the **All Profiles** list, expand the **AP** menu, then select **Dump collection**.
3. Configure the parameters described in Dump Collection Profile Configuration Parameters.
4. Click **Submit**.

5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check-box and then click **Deploy Changes**.

The following table describes the configuration parameters in the dump collection profile.

Table 136: *Dump Collection Profile Configuration Parameters*

Parameter	Description
Profile name	Enter a name for the dump collection profile.
Transfer Enable	Enable or disable APs to transfer core dump to controller and server. **NOTE:** By default, the `transfer-enable` sub-parameter is enabled allowing the AP to transfer the core dump and when disabled the core dump is stored on the AP.
Transfer Mode	Select the dump transfer mode from the drop down list: ■ TFTP ■ SCP ■ USB-Transfer
Server IP	Enter the server IP(v4/v6) for the dump collection.
Server Port	Enter the server port for the dump collection.
Server User name	Enter the server username for the dump collection.
Server Password	Enter the server password for the dump collection.
Server Path	Enter the server file path for the dump collection.
System Dump	Enable or disable the collection of system dump when an AP process crashes.

The following command configures the dump collection profile.

```
(host)  [mynode]  (config) #ap system-profile <profile-name>
(host) [mynode] (AP system profile "<profile-name>") #dump-collection-profile
```

Tri-Radio Mode for 550 Series Access Points

Starting from ArubaOS 8.6.0.0, 550 Series Access points will support 802.11ax 8x8 dual-radio with optional 4x4 tri-radio operating mode.

In tri-radio mode or split 5GHz mode, 8x8:8SS 5GHz radio is split into dual 4x4:4SS 5GHz radio. The two radios can work on AP mode and also work on AP+AM or AP+ Spectrum mode, where one radio provides wireless access and the other radio performs scanning. Tri-radio mode works only under BT POE or DC power.

The Tri-radio mode in 550 Series Access Points supports the following features:

- Station Management
- AirMatch
- SAPD/SAPM
- Spectrum Analysis
- Cluster

- MultiZone
- Mesh
- ClientMatch
- Firmware



When an AP is in a mode in which there are two radios on A-band,ClientMatch will not try to steer or load balance clients between the two A-band radios on the same AP. This limitation also applies to access points in dual-5G mode.

Follow the procedure to enable tri-radio mode in the WebUI:

1. In the **Managed Network** Node hierarchy, navigate to **Configuration > AP groups**.
2. Select an AP group. In the **AP group > <Name of the AP group> table**, select **Radios** and expand the **Advanced** accordion.
3. Under **5 GHz**, enable the **Split radio** toggle switch.
4. Enable the **Set second radio differently** toggle switch to select the radio mode.
5. Select the radio mode **am-mode / ap-mode / spectrum-mode** from the **Radio mode** drop-down list.
6. Click **Submit**.
7. Click **Pending Changes.**
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

For other technical specifications, refer 550 Series *Campus Access Points Installation Guide*.

AP USB Management

ArubaOS supports new infrastructure to manage any USB device that is plugged to an AP. The infrastructure allows describing a USB device through either CLI configuration or by using predefined descriptors. The infrastructure allows USB device management through USB ACLs. The infrastructure also supports plugin for USB devices. A plugin can perform further identification or perform normal logical interaction with the USB device. The infrastructure supports sending notification to other processes. For example: USB device plug-in or unplug notifications are sent to the IoT daemon. The infrastructure also supports sending script-based notifications. The USB ACLs are applied to the USB Management Daemon (UDMD) process and when a USB device is plugged in to an AP, the USB device walks through all USB ACL rules sequentially. ArubaOS supports USB devices only from Aruba approved vendors. The following vendors are approved:

- Alcatel-L800
- Amberbox-detector
- Amberbox-gateway
- C-motech-CNU-680
- EpiValley-SEC-8089
- Fraklin-u770-u772
- Franklin-U300
- Franklin-U301
- Franklin-U600
- Fujisoft

- Globetrotter-ICON-225
- Globetrotter-ICON-322
- HanShow
- Huawei-3276s-150
- Huawei-D41HW
- Huawei-E1552
- Huawei-E157
- Huawei-E160
- Huawei-E169-E180-E220
- Huawei-E170-E272-E220
- Huawei-E173
- Huawei-E1731-177DT06
- Huawei-E1750
- Huawei-E176-E176G-E1553
- Huawei-E1762
- Huawei-E180
- Huawei-E180-E1692-E1762
- Huawei-E1820e
- Huawei-E220
- Huawei-E261
- Huawei-E3131
- Huawei-E3272s-153
- Huawei-E3276
- Huawei-E3276s-500
- Huawei-E3372
- Huawei-E3372h-153-hil
- Huawei-E3372h-153-modem
- Huawei-E352s-5
- Huawei-E353
- Huawei-E353-E1750-E367
- Huawei-E367
- Huawei-E3765
- Huawei-E392
- Huawei-e398
- Huawei-E8372
- Huawei-EC150
- Huawei-EC167
- Huawei-HWD12-LTE
- Huawei-K3770
- Huawei-K3772
- Huawei-K4505
- Huawei-K4510
- Huawei-K4605

- Huawei-K5150
- Huawei-K5160
- Huawei-KDDI-DATA07
- Icon-452
- Longcheer-WM72
- Netgear-340u
- Netgear-341u
- Novatel-MC545
- Novatel-MC551L
- Novatel-MiFi-2200
- Novatel-Ovation-U727
- Novatel-U620L
- Novatel-U720
- Novatel-U727
- Novatel-U760-Sprint
- Novatel-U760-Virgin
- NTT-DoCoMo-L-02A
- NTT-DoCoMo-L-02C
- NTT-DoCoMo-L-05A
- NTT-DoCoMo-L-08C
- Pantech-UM150
- Pantech-UM175
- Pantech-UM190
- Pantech-UML290
- Pantech-UML295
- Pantech-UML295-cold
- Qualcomm-SXC-1080
- SES-Imagotag-021
- Sierra-250U
- Sierra-305-308
- Sierra-306-308-503-312U
- Sierra-313u
- Sierra-320U
- Sierra-330U
- Sierra-598
- Sierra-881U
- Sierra-885
- Sierra-Compass-597
- Sierra-Compass-885
- Sierra-Tstick-C597
- SIMTech
- Solu-M-SLG-DM101
- UGM1831

- UMG181
- Utstarcom-UM100C
- ZTE-3565
- ZTE-AC2726
- ZTE-AC2736
- ZTE-AC3781
- ZTE-Fivespot
- ZTE-K4505-z
- ZTE-MF110
- ZTE-MF180-HSDPA
- ZTE-MF190-Egypt
- ZTE-MF190-India
- ZTE-MF190-Thailand
- ZTE-MF591
- ZTE-MF633-MF636
- ZTE-MF637-MF656
- ZTE-MF668
- ZTE-MF683-HSDPA
- ZTE-MF79S
- ZTE-MF820
- ZTE-MF820D
- ZTE-MF823
- ZTE-MF825C
- ZTE-MF831
- ZTE-MF832S
- ZTE-MF832U
- ZTE-MF832U-Zero

USB ACL Profile

ArubaOS supports definition of up to 16 USB ACL profiles. A USB ACL profile is a vendor-product name with one action (permit, deny, reset). You cannot configure vendor ID in the USB descriptor, but can configure the name and action. When a USB ACL is applied to an AP group, all APs in that AP group inherit the USB ACL rules. When a USB device is described using a predefined descriptor, they are permitted by default.

When there is no USB ACL configuration but a USB profile is pushed, the UDMD process applies permit-all by default. If there is configured USB ACL, the SAPD or CLI process automatically adds deny-all rule at the end of the pushed rule list. That is, when UDMD does not find the matched USB ACL rule from the configuration, it applies the deny-all rule by default.

The following command creates a USB ACL profile named sample-usb-acl-profile with rule to permit USB devices from HanShow:

```
(host) (mynode) (config) #ap usb-acl-prof sample-usb-acl-profile
(host) (mynode) (AP USB ACL Profile "sample-usb-acl-profile") #rule vendor HanShow
action permit
```

The following command creates an AP USB profile named sample-ap-usb-profile and applies a USB ACL profile named sample-usb-acl-profile to it:

```
(host) (mynode) (config) #ap usb-profile sample-ap-usb-profile
(host) (mynode) (AP USB profile "sample-ap-usb-profile") #usb-acl-profile sample-usb-acl-profile
```

The following command applies a USB ACL profile named sample-usb-acl-profile to an AP group named sample-ap-group:

```
(host) (mynode) (config) #ap-group sample-ap-group
(host) (mynode) (config) #usb-profile sample-usb-acl-profile
```

The following command shows the USB profile:

```
(host) (mynode) #show ap usb-profile
```

The following command shows the details of a USB profile:

```
(host) (mynode) #show ap usb-profile sample-ap-usb-profile
```

The following command shows the USB ACL profile:

```
(host) (mynode) #show ap usb-acl-prof
```

The following command shows the details of a USB ACL profile:

```
(host) (mynode) #show ap usb-acl-prof sample-ap-usb-acl-profile
```

The following command shows the details of an AP group:

```
(host) (mynode) #show ap-group sample-ap-group
```

The following command shows the USB device management:

```
(host) (mynode) #show ap usb-device-mgmt [wired-mac <wired-mac>] | [ap-name <ap-name>] | [ip-addr <ip-addr>] | [ip6-addr <ip6-addr>] | all
```

The following command shows the debug USB device management:

```
(host) (mynode) #show ap debug usb-device-mgmt [ip-addr <ip-addr>] | [ip6-addr <ip6-addr>] | [ap-name <ap-name>] | [device all] | [module all] | [rule all] | [subscriber all] | [subscriber detail <service-type>] | [plugin all] | [plugin detail <device-id>]#
```

Disable AP Factory Reset

An AP may be reset to factory default configuration by pressing its reset button for more than 5 seconds while the AP is operational. ArubaOS allows to disable AP factory reset while the AP is operational.



By default, AP factory reset is enabled. That is, an AP may be reset to factory default configuration by pressing its reset button for more than 5 seconds while the AP is operational.

The following procedure describes how to disable AP factory reset while the AP is operational:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the **All Profiles** list, expand **AP** and select **AP system**.
3. Select the AP system profile you want to edit.
4. Expand **Advanced**.
5. Select the **Disable Factory Reset** check box.
6. Click **Submit**.

7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the required check box and click **Deploy Changes**.

The following CLI commands disable the AP factory reset while the AP is operational:

```
(host) [mynode] (config) #ap system-profile sample
(host) [mynode] (AP system profile "sample") #disable-factory-reset
```

The following procedure describes how to enable AP factory reset while the AP is operational:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the **All Profiles** list, expand **AP** and select **AP system**.
3. Select the AP system profile you want to edit.
4. Expand **Advanced**.
5. Unselect the **Disable Factory Reset** check box.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the required check box and click **Deploy Changes**.

The following CLI commands enable the AP factory reset while the AP is operational:

```
(host) [mynode] (config) #ap system-profile sample
(host) [mynode] (AP system profile "sample") #no disable-factory-reset
```

Chapter 28
Secure Enterprise Mesh

The Aruba 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. Aruba managed devices provide centralized configuration and management for APs in a mesh environment; local mesh APs provide encryption and traffic forwarding for mesh links.

The following topics in this chapter describe the components of the Aruba secure enterprise mesh architecture and profiles, as well as factors that should be taken into consideration when planning your mesh deployment.

- Overview of Mesh Access Points
- Overview of Mesh Links
- Overview of Mesh Profiles
- Overview of Remote Mesh Portals
- Overview of AP Boot Sequence
- Mesh Deployment Planning
- Mesh Deployment Solutions
- Mesh Configuration Procedures

Overview of Mesh Access Points

Mesh APs learn about their environment when they boot up. Mesh APs are either configured as a mesh portal, an AP that uses its wired interface to reach the managed device, or a mesh point, an AP that establishes an all-wireless path to the mesh portal. Mesh APs locate and associate with their nearest neighbor, which provides the best path to the mesh portal. Mesh portals and mesh points are also known as mesh nodes, a generic term used to describe APs configured for mesh.

A mesh radio's bandwidth can be shared between mesh-backhaul traffic and client traffic. You can, however, configure a radio for mesh services only. If you have a dual-radio AP, a mesh node can be configured to deliver client services on one radio, and both mesh and WLAN services to clients on the other. If you configure a single-radio AP to deliver mesh services only (by disabling the mesh radio in its 802.11a or 802.11g radio profile) that mesh node cannot deliver WLAN services to its clients.

For mesh and traditional thin AP deployments, the Aruba Mobility Master 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 WLAN links.

You configure the AP for mesh on the Mobility Master using either the WebUI or the CLI. All mesh related configuration parameters are grouped into mesh profiles that you can apply as needed to an AP group or to individual APs.

APs operate as thin APs by default; their primary function is to receive and transmit electromagnetic signals; other WLAN processing is left to the managed device. 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 managed device 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 managed device. This requires a minimum set of parameters from the AP group and mesh cluster so the mesh node discovers a neighbor, and creates a mesh link and subsequent channel with the managed device. 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 Configuring Ethernet Ports for Mesh or Provisioning Mesh Nodes.

The following sections provide information on the mesh portals, mesh points, and mesh clusters:

About Mesh Portals

The mesh portal is the gateway between the wireless mesh network and the enterprise wired LAN. You configure an Aruba 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 MSSID or 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 AES encryption. Mesh portals also propagate channel information, including CSAs.

About Mesh Points

The mesh point is an Aruba 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 Aruba WLAN services (such as client connectivity, IDS capabilities, user role association, LAN-to-LAN bridging, and QoS for LAN-to-mesh communication) to clients and performs mesh backhaul or 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 or untagged VLAN traffic across a mesh backhaul or network to a mesh portal.

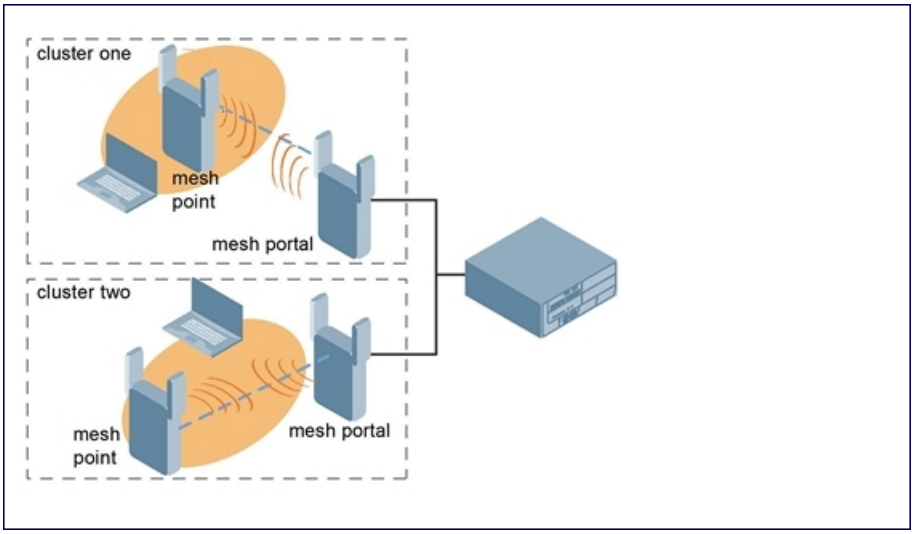
Mesh points use one of their wireless interfaces to carry traffic and reach the managed device. Mesh points are also aware of potential neighbors, and can form new mesh links if the current mesh link is no longer preferred or available.

About Mesh Clusters

Mesh clusters are similar to an ESS in a WLAN infrastructure. A mesh cluster is a logical set of mesh nodes that share the common connection and security parameters required to create mesh links. Mesh clusters are grouped and defined by a mesh cluster profile, as described in Mesh Cluster Profile.

Mesh clusters may enforce predictability in mesh networking by limiting the amount of concurrent mesh points, hop counts, and bandwidth used in the mesh network. A mesh cluster can have multiple mesh portals and mesh points that facilitate wireless communication between wired LANs. Mesh portals in a mesh cluster do not need to be on the same VLAN. Figure 88 shows two mesh clusters and their relationship to the managed device.

Figure 88 *Sample Mesh Clusters*



Mesh Network with Mixed Indoor and Outdoor APs

Indoor and outdoor APs participating in a mesh must be deployed with RF or regulatory settings, or assigned RF profiles, which allow overlapping channels to operate between the APs in the following scenarios:

- When setting up a mesh network between indoor and outdoor APs.
- When provisioning indoor APs as outdoor APs.
- When using AirMatch or ARM to support dynamic channel selection.

This allows both indoor and outdoor APs to always operate on the same channels. If you do not deploy the indoor and outdoor APs with regulatory settings or assigned RF profiles with overlapping channels, the indoor and outdoor channels for a given regulatory domain may not overlap. When the indoor and outdoor APs share a regulatory profile and are provisioned for the correct network environment, the dedicated indoor or outdoor mesh portals are deployed to support indoor or outdoor mesh points respectively.

The provisioned **Indoor** or **Outdoor** role of an AP is defined by the location of its antennas. Hence, when an indoor AP uses antennas installed in an outdoor area, the AP must be provisioned as **Outdoor**. For example, if the external outdoor antennas of an indoor AP are deployed to support outdoor APs as mesh points, and the indoor mesh portal is running on UNII-1 channel 36, then the outdoor mesh points may not be able to view the mesh portal to associate with. This occurs when the regulatory domain of that country has different allowed channels for indoor and outdoor APs, and the regulatory domain may disallow UNII-1 channels and UNII-3 channels for outdoor and indoor uses respectively. As a result, the mesh points cannot access UNII-1 APs. However, once the indoor AP with outdoor antennas is provisioned as an outdoor AP, that AP can then run on a UNII-3 channel, allowing the mesh points to access the portal.

Overview of Mesh Links

The mesh link is the data link between a mesh point and its parent. A mesh point uses the parameters defined in the mesh cluster profile, to establish a mesh link with a neighboring mesh point. The mesh link uses a series of metrics to establish the best path to the mesh portal.



The term **uplink** is used to distinguish the active association between a mesh point and its parent through this chapter.

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 MSSID from highest priority neighbors based on the least expected path cost.
- If no provisioned mesh cluster profile is available, mesh points use the recovery profile to establish an uplink. If multiple cluster profiles are configured, mesh points search, in order of priority, their list of provisioned backup mesh cluster profiles to establish an uplink. If the configured profiles are unavailable after searching for 5 minutes, the recovery profile is used.
- Moving to a better mesh link
- If the existing uplink quality degrades below the configured threshold, and a lower cost or more preferable uplink is available on the same channel and cluster, the mesh point reselects that link without re-scanning. In some cases, this invalidates all of the entries that have this mesh point as a next hop to the destination and triggers new learning of the bridge tables.
- Using a new mesh link if the current mesh link goes down
- If an uplink goes down, the affected mesh nodes 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.

The following sections provide information on link metrics and optimizing link metric algorithm:

About 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 137 describes the components that make up the metric value: node cost, hop count, link cost and 802.11 capacity.

The link metrics indicate the relative cost of a path to the mesh portal. The best path (lowest metric value) is used to create the uplink.

Table 137: *Mesh Link Metric Computation*

Component	Description
Node cost	Indicates the amount of traffic expected to traverse the mesh node. The more traffic, the higher the node cost. When establishing a mesh link, nodes with less traffic take precedence. The node cost is dependent on the number of children a mesh node supports. It can change as the mesh network topology changes, for example if new children are added to the network or old children disconnect from the network.
Hop count	Indicates the number of hops it takes the mesh node to get to the mesh portal. The mesh portal advertises a hop count of 0, while all other mesh nodes advertise a cumulative count based on the parent mesh node.

Component	Description
Link cost	Represents the quality of the link to an active neighbor. The higher the RSSI, the better the path to the neighbor and the mesh portal. If the RSSI value is below the configured threshold, the link cost is penalized to filter marginal links. A less direct, higher quality link may be preferred over the marginal link. The following factors also affect mesh link metrics: High-throughput APs add a high cost penalty for links to non-high-throughput APs. Multi-stream high-through APs add proportional cost penalties for links to high-throughput APs that support fewer streams.
802.11 capacity	High-throughput APs can send 802.11 information elements in their management frames, allowing high-throughput mesh nodes to identify other mesh nodes with a high-throughput capacity. High-throughput mesh points prefer to select other 802.11-capable mesh points in their path to the mesh portal, but can use a legacy path if no high-throughput path is available.
Path Cost	Path cost is calculated by analyzing the other components in this table, and adding the link cost, the mesh parent's path cost, and the parent's node cost. Mesh portals typically advertise a path-cost of zero, but high-throughput portals add an offset penalty if they are connected to a 10/100 mbps port that is too slow for the high-throughput link capacity.

About Optimizing Links

You can configure and optimize operation of the link metric algorithm through the mesh radio profile. These configurable mesh link trigger thresholds can determine when the uplink or mesh path is dropped and another is chosen, provide enhanced network reliability, and contain flapping links. Although you can modify the behavior of the link metric algorithm, It is recommended to follow the default values for most deployments.

Overview of 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 Mobility Master; you do not explicitly configure the recovery profile.

Aruba 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.

The various types of mesh profiles are described below:

About Mesh Cluster Profiles

Mesh clusters are grouped and defined by a mesh cluster profile, which provides the framework of the mesh network. Similar to virtual AP profiles, the mesh cluster profile contains the MSSID, authentication methods, security credentials, and cluster priority required for mesh nodes to associate with their

neighbors and join the cluster. Associated mesh nodes store this information in flash memory. Although most mesh deployments require only a single mesh cluster profile, you can configure and apply multiple mesh cluster profiles to an AP group or an individual AP. If you have multiple cluster profiles, the mesh portal uses the profile with the highest priority to bring up the mesh network. Mesh points, in contrast, go through the list of mesh cluster profiles in order of priority to decide which profile to use to associate themselves with the network. The mesh cluster priority determines the order by which the mesh cluster profiles are used. This allows you, rather than the link metric algorithm, to explicitly segment the network by defining multiple cluster profiles.

Since the mesh cluster profile provides the framework of the mesh network, you must define and configure the mesh cluster profile before configuring an AP to operate as a mesh node. You can use either the default cluster profile or create your own. If you find it necessary to define more than one mesh cluster profile, you must assign priorities to each profile to allow the Mesh AP group to identify the primary and backup mesh cluster profiles. 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 access point, this prevents a mesh cluster profile defined at the AP group level from being applied to a specific AP.

Do not delete or modify mesh cluster profiles once you use them to provision mesh nodes. You can recover the mesh point if the original cluster profile is still available. It is recommended to create a new mesh cluster profile if needed. If you modify any mesh cluster setting, you must reprovision your AP for the changes to take effect (this also causes the AP to automatically reboot).

If you configure multiple cluster profiles with different cluster priorities, you manually override the link metric algorithm because the priority takes precedence over the path cost. In this scenario, the mesh portal uses the profile with the highest priority to bring-up the mesh network. The mesh portal stores and advertises that one profile to neighboring mesh nodes to build the mesh network. This profile is known as the primary cluster profile. Mesh points, in contrast, go through the list of configured mesh cluster profiles in order of priority to find the profile being advertised by the mesh portal. Once the primary profile has been identified, the other profiles are considered backup cluster profiles. Use this deployment if you want to enforce a particular mesh topology rather than allowing the link metric algorithm to determine the topology.

For this scenario, do the following:

- Configure multiple mesh cluster profiles with different priorities. The primary cluster profile has a lower priority number, which gives it a higher priority.
- Configure the mesh radio profile.
- Create an AP group for 802.11a radios and 802.11g radios
- Configure the 802.11a or 802.11g RF management profiles for each AP group.
- If your deployment includes high-throughput APs, configure the mesh high-throughput SSID profile. The mesh radio profile uses the default high-throughput SSID profile unless you specifically configure the mesh radio profile to use a different high-throughput SSID profile
- Create an AP group for each 802.11a channel.

If a mesh link breaks or the primary cluster profile is unavailable, mesh nodes use the highest priority backup cluster profile to re-establish the uplink or check for parents in the backup profiles. If these profiles are unavailable, the mesh node can revert to the recovery profile to bring up the mesh network until a cluster profile is available.

About Mesh Radio Profiles

The mesh radio profile allows you to specify the set of rates used to transmit data on the mesh link. This profile also allows you to define a **reselection-mode** setting to optimize the operation of the link metric algorithm. The reselection mode specifies the method a mesh node uses to find a better uplink to create a path to the mesh portal. Only neighbors on the same channel in the same mesh cluster are considered.

The mesh radio profile includes the following reselection mode options:

- **reselect- anytime**: mesh points using the **reselect-anytime** reselection mode perform a single topology readjustment scan within 9 minutes of startup and 4 minutes after a link is formed. If no better parent is found, the mesh point returns to its original parent. This initial scan evaluates more distant mesh points before closer mesh points, and incurs a dropout of 5–8 seconds for each mesh point. After the initial startup scan is completed, connected mesh nodes evaluate mesh links every 30 seconds. If a mesh node finds a better uplink, the mesh node connects to the new parent to create an improved path to the mesh portal.
- **reselect-never**: connected mesh nodes do not evaluate other mesh links to create an improved path to the mesh portal.
- **startup-subthreshold**: mesh points using the **startup-subthreshold** reselection mode perform a single topology readjustment scan within 9 minutes of startup and 4 minutes after a link is formed. If no better parent is found, the mesh point returns to its original parent. This initial startup scan evaluates more distant mesh points before closer mesh points, and incurs a dropout of 5–8 seconds for each mesh point. After that time, each mesh node evaluates alternative links if the existing uplink falls below the configured threshold level (the link becomes a sub-threshold link). It is recommended to use this default **startup-subthreshold** value.
- **subthreshold-only**: connected mesh nodes evaluate alternative links only if the existing uplink becomes a sub-threshold link.

If a mesh point using the **startup-subthreshold** or **subthreshold-only** mode reselects a more distant parent because its original, closer parent falls below the acceptable threshold, then as long as that mesh point is connected to that more distant parent, it seeks to reselect a parent at the earlier, shorter distance (or less) with good link quality. For example, if a mesh point disconnects from a mesh parent 2 hops away and subsequently reconnects to a mesh parent 3 hops away, then the mesh point continues to seek a connection to a mesh parent with both an acceptable link quality and a distance of two hops or less, even if the more distant parent also has an acceptable link quality.

About 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 802.11g (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 2.4 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.



If you do no want the mesh radios carrying mesh-backhaul traffic to support client traffic, consider using a dedicated 802.11a or 802.11g radio profile with the mesh radio disabled. In this scenario, the radio carries mesh backhaul traffic but does not support client virtual APs.

Mesh nodes operating in different cluster profiles can share the same radio profile. Conversely, mesh portals using the same cluster profile can be assigned different RF Management Radio profiles to achieve frequency separation.

About High-Throughput Radio Profiles

Each 802.11a and 802.11g radio profile also references a high-throughput profile that manages an AP or AP group's 40 MHz tolerance settings.

About Mesh High-Throughput SSID Profiles

High-throughput APs support additional settings not available in legacy APs. A mesh high-throughput SSID profile can enable or disable high-throughput (802.11n) features and 40 MHz channel usage, and define values A-MDPUs and MCS ranges.

Aruba 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.

About Wired AP Profiles

The wired AP profile controls the configuration of the Ethernet ports 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.

About Mesh Recovery Profiles

In addition to the default and user-defined mesh cluster profiles, mesh nodes also have a recovery profile. The Mobility Master dynamically generates a recovery profile, and each mesh node provisioned by the same Mobility Master has the same recovery profile. The recovery profile is based on a PSK, and mesh nodes use the recovery profile to establish a link to the managed device if the mesh link is broken and no other mesh cluster profiles are available.

The mesh portal advertises the provisioned cluster profile. If a mesh point is unaware of the active mesh cluster profile, but is aware of and has the same recovery profile as the mesh portal, the mesh point can use the recovery profile to connect to the mesh portal.



The mesh point must have the same recovery profile as the parent to which it connects. If you provision the mesh points with the same Mobility Master, 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 managed device, 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 Mobility Master's configuration file and is unique to that Mobility Master. If necessary, you can transfer your configuration to another managed device. If you do so, make sure your new mesh cluster is running and you have re-provisioned the mesh nodes before deleting your previous configuration. The APs learn the new recovery profile after they are provisioned with the new managed device. This is also true if you provision a mesh node with one Mobility Master and use it

with a different Mobility Master. In this case, the recovery profile does not work on the mesh node until you re-provision it with the new Mobility Master.

Overview of Remote Mesh Portals

You can deploy mesh portals to create a hybrid mesh or remote AP environment to extend network coverage to remote locations; this feature is called Remote Mesh Portal. The Remote Mesh Portal feature integrates the functions of a Remote AP and the Mesh portal. As a Remote AP, it sets up a VPN tunnel back to the corporate switch that secures control traffic between the Remote AP 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 Mobility Master. 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.

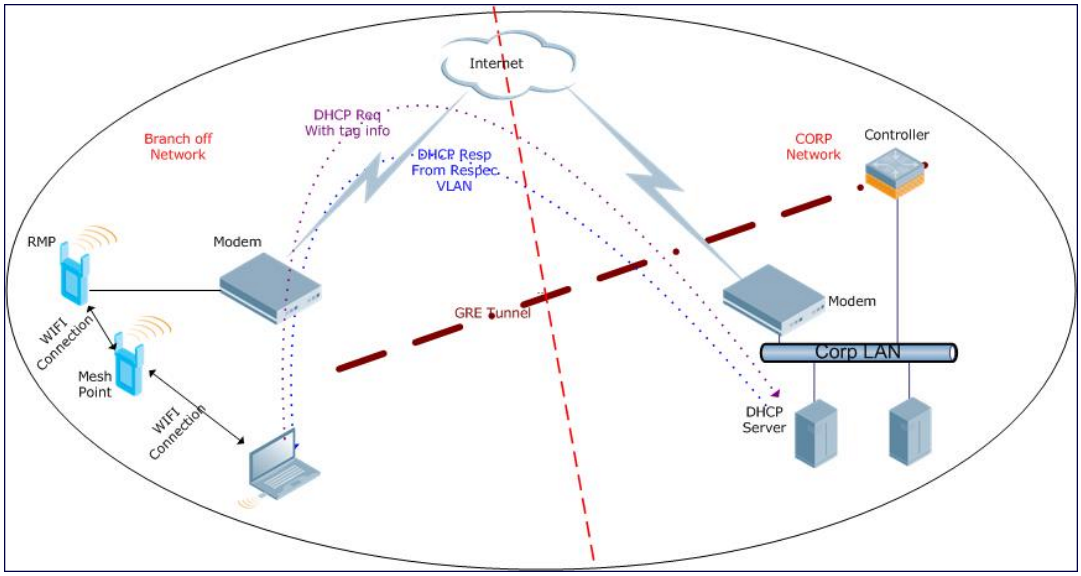
When the client at the branch office associates to a virtual AP in split-tunnel forwarding mode, the client's DHCP requests are forwarded over a GRE tunnel (split tunnel) to the corporate network. This communication is done over a secure VPN tunnel. The IPs are assigned from the corporate pool based 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 mesh private VLAN parameter. The mesh point learns the mesh private VLAN value from the response during the mesh association. When the split tunnel is set up for the remote mesh portal on the Mobility Master, the VLAN of the tunnel should be the mesh private VLAN. A DHCP pool for the mesh private VLAN should be set up on the switch. The use of mesh private VLAN makes it easy for the remote mesh portal to decide which requests to forward over the split tunnel. All requests tagged with the mesh private VLAN are sent over the split tunnel. Hence, the mesh private VLAN should be different from any user VLAN that is bridged using the mesh network.



The remote mesh portal configuration requires an AP license. For complete information on the licenses, refer to the *Aruba Mobility Master Licensing Guide.*

Figure 89 *Working of Remote Mesh Portal*



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 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 managed device using these IP addresses. When these mesh points send and receive PAPI control traffic from the Mobility Master, 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 managed device 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 through 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 mesh private VLAN value so that it does not clash with any local tags assigned in the mesh network. In this scenario, the portal performs the default operation and bridges the frame based on its bridge table. Traffic destined to the Internet is recognized as such by the remote mesh portal based on ACL rules. This traffic is NATed on the remote mesh portal's Ethernet interface.

Overview of AP Boot Sequence

This section describes the boot sequence for mesh APs in detail. Depending on its configured role, the AP performs a slightly different boot sequence.

Booting the Mesh Portal

When the mesh portal boots, it recognizes that one radio is configured to operate as a mesh portal. It then obtains an IP address from a DHCP server on its Ethernet interface, discovers the Mobility Master on that interface, registers the mesh radio with the managed device, 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 set up the mesh virtual AP with the correct announcements on beacons and probe responses. On the non-mesh radio provisioned for access mode, that radio is a thin AP and everything on that interface works as a thin AP radio interface.

If the 802.11a or 802.11g radio profile assigned to the mesh radio is enabled, the radio supports both mesh backhaul and client access Virtual APs. If the mesh radio is to be used exclusively for mesh backhaul traffic, associate that radio to a dedicated 802.11a or 802.11g radio profile with the radio disabled so the mesh radios carry backhaul traffic only.

Booting the Mesh Point

When the mesh point boots, it scans for neighboring mesh nodes to establish a link to the mesh portal. All of the mesh nodes that establish the link are in the same mesh cluster. After the link is up, the mesh point uses the DHCP to obtain an IP address and uses the same Mobility Master as their parent. The remaining boot sequence, if applicable, is similar to that of a thin AP. Remember, the priority of the mesh point is establishing a link with neighboring mesh nodes, not establishing a control link to a managed device.



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 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 on the managed node and AMs deployed in your network to distinguish between APs, wireless clients, and mesh nodes. The WMS tables also identify the mesh nodes.

For all thin APs and mesh nodes, the AM identifies a mesh node from other packets monitored on the air, and the AM does not trigger wireless-bridging events for packets transmitted between mesh nodes.

Mesh Deployment Planning

Following considerations are recommended when planning and deploying a mesh solution:

Pre-Deployment Considerations

- Stage the APs before deployment. Identify the location of the APs, configure them for mesh, provision them, and verify connectivity before physically deploying the mesh APs in a live network.
- Ensure the Mobility Master has Layer-2 or Layer-3 network connectivity to the network segment where you plan to install the mesh portal.
- Keep the AP packaging materials and reuse them to send the APs to the installation location.
- Verify the layout of the physical location to determine the appropriate configuration and placement of the APs. Use this information to avoid problems that would necessitate a physical recovery.
- Label the AP before sending it to the physical location for installation.

Outdoor-Specific Deployment Considerations

- Provision the AP with the latitude and longitude coordinates of the installation location. This allows you to more easily identify the AP for inventory and troubleshooting purposes.
- Identify a radio line of sight between the antennas for optimum performance. The radio line of sight involves the area along a link through which the bulk of the radio signal power travels.
- Identify the minimum antenna height required to ensure a reliable mesh link.
- Scan your proposed site to avoid radio interference caused by other radio transmissions using the same or an adjacent frequency.
- Consider extreme weather conditions known to affect your location, including: temperature, wind velocity, lightning, rain, snow, and ice.
- Allow for seasonal variations, such as growth of foliage.

For more detailed outdoor deployment information, refer to the installation guide that came with your outdoor AP.

Configuration Considerations

- On dual-radio APs, you can configure only one of the radio for mesh. If you want a dual-radio AP to carry mesh backhaul traffic and client services traffic on separate radios, it is recommended to use 802.11a radios for mesh-backhaul traffic and 802.11g radios for traditional WLAN access.
- If you configure more than one mesh node in the same VLAN, prevent network loops by enabling STP on the Layer-2 switch used to connect the mesh nodes.
- Mesh nodes learn a maximum of 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.
- APs configured as mesh points support secure jack operation on Ethernet 0. APs with multiple Ethernet ports configured as mesh portals support secure jack operation on Ethernet 1. If an AP with multiple Ethernet ports is configured as a mesh point, it supports secure jack operation on Ethernet 1 and Ethernet 0.
- Mesh networks forward tagged or untagged VLAN traffic, but do not tag traffic. The allowed VLANS are controlled by the wired ap profile.
- Mesh APs provisioned on different managed device can interoperate if those APs are configured with the same country code, cluster name and cluster key. However, the mesh recovery profile created on one managed device is not able to recover settings for mesh APs provisioned on another managed device unless the recovery profile is on Mobility Master and the other mesh nodes were provisioned by a managed device connected to that master.

Post-Deployment Considerations

- Do not connect mesh point Ethernet ports in such a way that causes a network loop.
- Have a trained professional install the AP. After installation, check to ensure the AP receives power and boots up, enabling RSSI outputs.



Although the AP is up and operational, it is not connected to the network.

- Align the AP antenna for optimal RSSI.
- Do not delete or modify mesh cluster profiles once you use them to provision mesh nodes. You can recover the mesh point if the original cluster profile is still available. It is recommended to create a new mesh cluster profile if needed.
- If you create a new mesh cluster profile for an existing deployment, you must re-provision the AP for the new profile to take effect. If you re-provision mesh nodes that are already operating, re-provision the most distant (highest hop count) mesh points first, followed by the mesh portals. If you re-provision the mesh portal first, the mesh points may be unable to form a mesh link. Note that re-provisioning the AP causes it to automatically reboot, which may cause a disruption of service to the network.

Dual-Port AP Considerations

A dual-port AP has two 10/100 Mbps Ethernet ports (Ethernet 0 and Ethernet 1, respectively). When using these APs in a mesh environment, note the following Ethernet port requirements:

- If configured as a mesh portal:
  - Connect Ethernet 0 to the managed device to obtain an IP address. The wired AP profile controls Ethernet 1.
  - Only Ethernet 1 supports secure jack operation.
- If configured as a mesh point, Ethernet 0 and Ethernet 1 can be configured using separate wired-port-profiles.

Mesh Deployment Solutions

- You can configure the following single-hop and multi-hop solutions:
- Thin AP services with wireless backhaul deployment
- Point-to-point deployment

- Point-to-multipoint deployment
- High-availability deployment

With a thin AP wireless backhaul deployment, mesh provides services and security to remote wireless clients and sends all control and user traffic to the Mobility Master over a wireless backhaul mesh link.

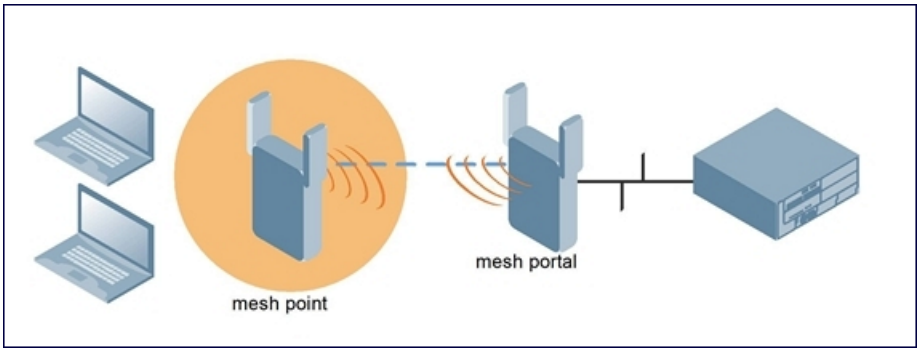
The remaining deployments allow you to extend your existing wired network by providing a wireless bridge to connect Ethernet LAN segments. You can use these deployments to bridge Ethernet LANs between floors, office buildings, campuses, factories, warehouses, and other environments where you do not have access to physical ports, or cable to extend the wired network. In these scenarios, a wireless backhaul carries traffic between the Aruba APs configured as the mesh portal and the mesh point, to the Ethernet LAN.

The following topics provide information on the single-hop and multi-hop mesh deployment solutions:

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 managed device. Use the 802.11g radio for WLAN and managed device services and the 802.11a radio for mesh services. Figure 90 shows the wireless backhaul between the mesh portal to the mesh point that services the wireless clients.

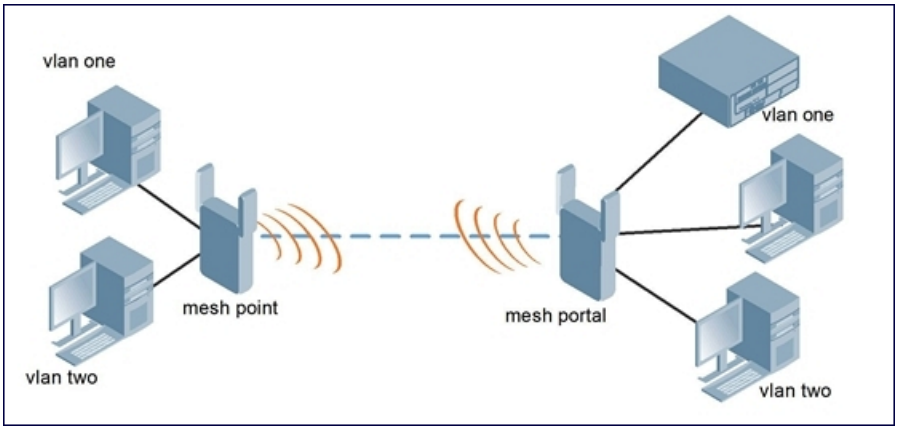
Figure 90 *Sample Wireless Backhaul Deployment*



Point-to-Point Deployment

In this point-to-point scenario, two Ethernet LAN segments are bridged through 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 91 shows a single-hop point-to-point deployment.

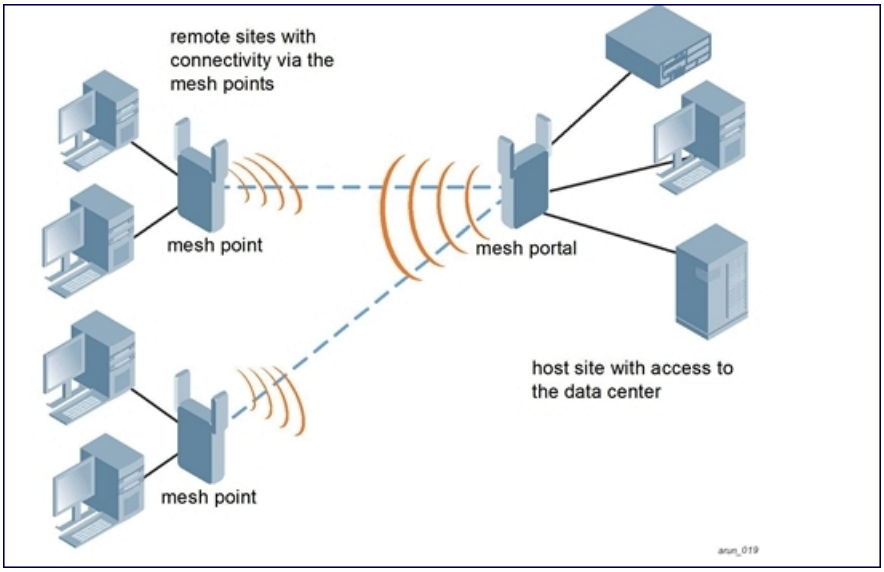
Figure 91 *Sample Point-to-Point Deployment*



Point-to-Multipoint Deployment

In a point-to-multipoint scenario, multiple Ethernet LAN segments are bridged through multiple wireless or 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 92 shows a single-hop point-to-multipoint deployment.

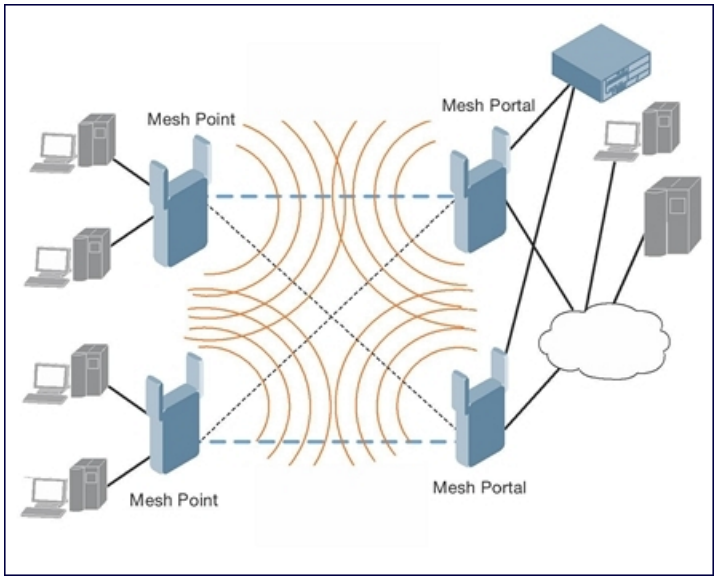
Figure 92 *Sample Point-to-Multipoint Deployment*



High-Availability Deployment

In this high-availability scenario, multiple Ethernet LAN segments are bridged through 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 93 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 93 *Sample High-Availability Deployment*



Mesh Configuration Procedures

The following topics describe the procedures required to configure your secure enterprise mesh solution:

1. Creating and Editing Mesh Radio Profiles
2. Creating and Editing Mesh High-Throughput SSID Profiles
3. Configuring Mesh Cluster Profiles
4. Configuring Mesh Clusters Associated with AP Groups
5. Configuring Ethernet Ports for Mesh
6. Configuring a Mesh Access List
7. Provisioning Mesh Nodes
8. Radio Selection for Mesh Links
9. Verifying Your Mesh Network

NOTE

Aruba strongly recommends staging mesh APs before deploying them. Identify the physical location of the APs, configure them for mesh, provision the APs and verify connectivity before physically deploying them in a live network.

If you are configuring an AP as both a remote access point and a mesh portal, see also Configuring Remote Mesh Portals

Creating and Editing Mesh Radio Profiles

The mesh radio profile determines many of the settings used by mesh nodes to establish mesh links and the path to the mesh portal, including the maximum number of children a mesh node can accept, and transmit rates for the 802.11a and 802.11g radios. The attributes of the mesh radio profile are applied to a mesh point upon receiving its configuration from the managed node. You can configure multiple radio profiles; however, you select and deploy only one radio profile per AP group. Radio profiles, including the default profile, are not active until you provision your APs for mesh.

If you modify a currently provisioned and running radio profile, your changes take effect immediately. You do not need to reboot the managed device or the AP to apply the changes.

This section describes the following topics:

- Managing Mesh Radio Profiles in the WebUI
- Managing Mesh Radio Profiles in the CLI

Managing Mesh Radio Profiles in the WebUI

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

- Creating or Editing a Mesh Radio Profile
- Assigning a Mesh Radio Profile to an AP Group

Creating or Editing a Mesh Radio Profile

The following procedure describes how to create or edit an existing mesh radio profile:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Under the **All Profiles** pane, expand **Mesh** and select **Mesh Radio**.
3. The procedure to create a new mesh profile varies slightly from the procedure to edit an existing profile.
   - To create a new mesh profile, click **+** in the **Mesh Radio Profile: New Profile** and enter the **Profile Name.**
   - To edit an existing mesh profile, select the profile that you want to edit from **Mesh > Mesh Radio**.
4. Configure your desired mesh radio settings.
5. Mesh Radio profile configuration settings are divided into two tabs, **General** and **Advanced**. The **General** tab displays only those configuration settings that often need to be adjusted to suit a specific network. The **Advanced** tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values. If you change a setting on one tab, then click and display the other tab without saving your configuration, that setting reverts to its previous value. The basic and advanced profile settings are described in Table 138.

Table 138: *Mesh Radio Profile Configuration Parameters*

Parameter	Description
General Mesh Radio Settings	
Link Threshold	Use this setting to optimize operation of the link metric algorithm. Indicates the minimal RSSI value. If the RSSI value is below this threshold, the link may be considered a sub-threshold link. A sub-threshold link is one whose average RSSI value falls below the configured link threshold. If this occurs, the mesh node may try to find a better link on the same channel and cluster (only neighbors on the same channel are considered). Default: 12. The supported threshold is hardware dependent, with a practical range of 10–90.
Advanced Mesh Radio Settings	

Parameter	Description
802.11a Transmit Rates	Indicates the transmit rates for the 802.11a radio. The AP attempts to use the highest transmission rate to establish a mesh link. If a rate is unavailable, the AP goes through the list and uses the next highest rate. To modify transmit rates, do one of the following: In the WebUI, deselect (uncheck) a specific rate box to use fewer rates when establishing a mesh link. In the CLI, enter the specific rates to use. Default: All transmission rates are selected and used. If you do not select 802.11a or 802.11g transmit rates, all rates are selected by default when you click **Save**.
802.11g Transmit Rates	Indicates the transmit rates for the 802.11g radio. The AP attempts to use the highest transmission rate to establish a mesh link. If a rate is unavailable, the AP goes through the list and uses the next highest rate. To modify transmit rates, do one of the following: In the WebUI, deselect (uncheck) a specific rate box to use fewer rates when establishing a mesh link. In the CLI, enter the specific rates to use. Default: All transmission rates are selected and used. If you do not select 802.11a or 802.11g transmit rates, all rates are selected by default when you click Save.
Allowed VLANs on Mesh Link	List the VLAN ID numbers of VLANs allowed on the mesh link.
BC/MC Rate Optimization	Broadcast or Multicast Rate Optimization dynamically selects the rate for sending broadcast or multicast frames on any BSS. This feature determines the optimal rate for sending broadcast and multicast frames based on the lowest of the unicast rates across all associated clients. When you enable the Multicast Rate Optimization feature, the managed node scans the list of all associated stations in that BSS and finds the lowest transmission rate as indicated by the rate adaptation state for each station. If there are no associated stations in the BSS, it selects the lowest configured rate as the transmission rate for broadcast and multicast frames. This feature is enabled by default. Multicast Rate Optimization applies to broadcast and multicast frames only. 802.11 management frames are not affected by this feature and are transmitted at the lowest configured rate. When enabled, this setting dynamically adjusts the multicast rate to that of the slowest connected mesh child. Multicast frames are not sent if there are no mesh children. 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 or multicast packets at that station. Configuring the **Video Multicast Rate Optimization** parameter overrides the configuration of **BC/MC Rate Optimization** parameter for VI-tagged multicast traffic. Multicast traffic that is not VI-tagged behaves the same with BC or MC as before. If multicast rate is not set, all traffic behaves the same. Default: Enabled.

Parameter	Description
Heartbeat Threshold	Indicates the maximum number of heartbeat messages that can be lost between neighboring mesh nodes. Default: 10 missed heartbeats. Range: 1–255.
Maximum Children	Indicates the maximum number of children a mesh node can accept. Default: 64 children. Range: 1–64
Maximum Hop Count	Indicates the maximum hop count from the mesh portal. Default: 8 hops. Range: 1–32
Mesh Private VLAN	A VLAN ID for control traffic between an remote mesh portal and mesh nodes. This VLAN ID must not be used for user traffic. Range: 0–4094. Default: 0 (disabled).
Metric algorithm	This parameter specifies the algorithm used by a mesh node to select its parent. Use this setting to optimize operation of the link metric algorithm. Available options are: best-link-rssi: Selects the parent with the strongest RSSI, regardless of the number of children a potential parent has. distributed-tree-rssi: selects the parent based on link-RSSI and node cost based on the number of children. This option evenly distributes the mesh points over high quality uplinks. Low quality uplinks are selected as a last resort. Default: distributed-tree-rssi. It is recommended to use the default value.
Rate Optimization for delivering EAPOL frames and mesh echoes	When you enable this parameter, EAPOL frames, mesh echo requests and echo responses are sent at a lower rate.
Reselection Mode	Use this setting to optimize operation of the link metric algorithm. Available options are: reselect-anytime reselect-never startup-subthreshold subthreshold-only
Retry Limit	Indicates the number of times a mesh node can re-send a packet. Default: 4 times. Range: 1–15
RTS Threshold	Defines the packet size sent by mesh nodes. Mesh nodes transmitting frames larger than this threshold must issue RTS and wait for other mesh nodes to respond with CTS to begin transmission. This helps prevent mid-air collisions. Default: 2333 bytes. Range: 256– 2346.

6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Assigning a Mesh Radio Profile to an AP Group

The following procedure describes how to associate a mesh radio profile to a mesh AP or AP:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > AP Groups**.
2. Select an AP group.
3. Click the **Profiles** tab.
4. Under the **Profiles** list, expand **Mesh**, and then select **Mesh Radio**.
5. Select a profile from the **Mesh Radio profile** drop-down list.
6. Open the **Mesh High-throughput SSID** configuration for the radio profile and select an SSID profile from the **Mesh High-throughput SSID profile** drop-down list.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

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.

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

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.

Managing Mesh Radio Profiles in the CLI

You must be in configuration 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.

This section describes the following topics:

- Creating or Modifying a Mesh Radio Profile
- Assigning a Mesh Radio Profile to a Mesh AP or AP Group
- Deleting a Mesh Radio Profile

Creating or Modifying a Mesh Radio Profile

Configuration details and any default values for each of these parameters are described in . 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.

The following CLI commands create or modify mesh radio profiles:

```
(host) [mynode] (config) #ap mesh-radio-profile <profile-name>
a-tx-rates
allowed-vlans
children <children>
clone <source-profile-name>
eapol-rate-opt
g-tx-rates [1|2|5|6|9|11|12|18|24|36|48|54]
heartbeat-threshold <count>
hop-count <hop-count>
link-threshold <count>
max-retries <max-retries>
mesh-ht-ssid-profile
mesh-mcast-opt
mesh-survivability
```

```
metric-algorithm {best-link-rssi|distributed-tree-rssi}
mpv <vlan-id>
no
reselection-mode
rts-threshold <rts-threshold>
```

You can also create a new mesh radio profile by copying the settings of an existing profile using the **clone** parameter. Using the **clone** command to create a new profile makes it easier to keep constant attributes in common within multiple profiles.

```
(host) [mynode] (config) #ap mesh-radio-profile <profile-name> clone <source-profile-
name>
```

Assigning a Mesh Radio Profile to a Mesh AP or AP Group

The following CLI commands associate a mesh radio profile with an AP or AP group. When you add the mesh cluster profile to the AP group, you must also define the cluster priority.

```
(host) [mynode] (config) #ap-group <group>
mesh-radio-profile <profile-name> priority <priority>
```

The following CLI commands associate a mesh radio profile with an individual AP:

```
(host) [mynode] (config) #ap-name <name>
mesh-radio-profile <profile-name> priority <priority>
```

The following CLI command assigns the mesh cluster profiles **cluster1** and **cluster2** to two different AP groups. In the AP group **group1, cluster1** has a priority of 5, and **cluster2** has a priority of 10, so **cluster1** has the higher priority. In the AP group **group2**, **cluster1** has a priority of 10, and **cluster2** has a priority of 5, so **cluster5** has the higher priority.

```
(host) [mynode] (config) #ap-group group1
   mesh-cluster-profile cluster1 priority 5
   mesh-cluster-profile cluster2 priority 10

(host) [mynode] (config) #ap-group group2
mesh-cluster-profile cluster1 priority 10
mesh-cluster-profile cluster2 priority 5
```

Deleting a Mesh Radio Profile

The following CLI command deletes a radio profile via the command-line interface.

```
(host) [mynode] (config) #no ap mesh-radio-profile <profile-name>
```

Creating and Editing Mesh High-Throughput SSID Profiles

The mesh high-throughput SSID profile defines settings unique to 802.11n and 802.11ac-capable, high-throughput APs. If none of the APs in your mesh deployment are 802.11n or 802.11ac-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 need to reboot the managed device or the AP.

This section describes the following topics:

- Managing Mesh High-Throughput SSID Profiles in the WebUI
- Managing Mesh High-Throughput SSID Profiles in the CLI

Managing Mesh High-Throughput SSID Profiles in the WebUI

Use the following procedures to manage your high-throughput SSID profiles.

This section contains the following topics:

1. Creating or Editing a Profile
2. Assigning a Profile to an AP Group
3. Deleting a Profile

Creating or Editing a Profile

The following procedure describes how to create a high-throughput SSID profile:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System>Profiles** tab.
2. Under the **All Profiles** menu, expand the **Mesh** tab and select the **Mesh High-throughput SSID** profile.
3. To add a new profile, click **+** in **Mesh High-throughput SSID profile: New Profile.**
4. Enter a **Profile Name**.
5. Configure the mesh high-throughput SSID parameters described in Table 139. The Mesh High-Throughput SSID Profile configuration settings are divided into three tabs, **General**, **Transmit Beamforming**, and **Advanced**. The tab displays only those configuration settings that often need to be adjusted to suit a specific network. The tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values. If you change a setting on one tab then click and display the other tab without saving your configuration, that setting reverts to its previous value.

Table 139: *Mesh High-Throughput SSID Profile Configuration Parameters*

Parameter	Description
High-throughput Enable (SSID)	Enable or disable high-throughput (802.11n) features on the SSID. Default: enabled
40 MHz channel usage	Enable or disable the use of 40 MHz channels. Default: enabled
Very High throughput enable (VHT)	Enable or disable very high-throughput (802.11av) features on the SSID. Default: enabled
80 MHz channel usage	Enable or disable the use of 80 MHz channels. Default: enabled
VHT- Explicit Transmit Beamforming	Enable or disable use of Very High Through-put Explicit Transmit Beamforming. If this parameter is disabled, the other transmit beamforming configuration settings have no effect.
BA AMSDU Enable	Enable or disable Receive AMSDU in BA negotiation.
Temporal Diversity Enable	When a client is not responding to 802.11 packets, the AP will launch two hardware retries. If you enable this option and hardware retries are not successful, then the AP will launch and the software retries.
Legacy stations	Allow or disallow associations from legacy (non-HT) stations. By default, this parameter is enabled (legacy stations are allowed).

Parameter	Description
Low-density Parity Check	If enabled, the AP advertises LDPC support. LDPC improves data transmission over radio channels with high levels of background noise.
MPDU Aggregation	Enable or disable MPDU aggregation. High-throughput APs are able to send aggregated MAC protocol data units MDPUs, which allow an AP to receive a single block acknowledgment instead of multiple ACK signals. This option, which is enabled by default, reduces network traffic overhead by effectively eliminating the need to initiate a new transfer for every MPDU.
Max received A-MPDU size	Maximum size of a received aggregate MPDU, in bytes. Allowed values: 8191, 16383, 32767, 65535.
Max transmitted A-MPDU size	Maximum size of a transmitted aggregate MPDU, in bytes. Range: 1576–65535
Min MPDU start spacing	Minimum time between the start of adjacent MPDUs within an aggregate MPDU, in microseconds. Allowed values: 0 (No restriction on MDPU start spacing), .25 µsec, .5 µsec, 1 µsec, 2 µsec, 4 µsec.
Short guard interval in 20 MHz mode	Enable or disable use of short (400 ns) guard interval in 20 MHz mode. This parameter is enabled by default. A guard interval is a period of time between transmissions that allows reflections from the previous data transmission to settle before an AP transmits data again. An AP identifies any signal content received inside this interval as unwanted inter-symbol interference, and rejects that data. The 802.11n standard specifies two guard intervals: 400 ns (short) and 800 ns (long). Enabling a short guard interval can decrease network overhead by reducing unnecessary idle time on each AP. Some outdoor deployments, may, however require a longer guard interval. If the short guard interval does not allow enough time for reflections to settle in your mesh deployment, inter-symbol interference values may increase and degrade throughput.
Short guard interval in 40 MHz mode	Enable or disable use of short (400 ns) 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: 400 ns (short) and 800 ns (long). Enabling a short guard interval can decrease network overhead by reducing unnecessary idle time on each AP. Some outdoor deployments, may, however require a longer guard interval. If the short guard interval does not allow enough time for reflections to settle in your mesh deployment, inter-symbol interference values may increase and degrade throughput.
Short guard interval in 80 MHz mode	Enable or disable use of short (400 ns) guard interval in 80 MHz mode. 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.

Parameter	Description
	The 802.11n standard specifies two guard intervals: 400 ns (short) and 800 ns (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. This parameter is enabled by default.
Supported MCS set	A list of MCS values or ranges of values to be supported on this SSID. The MCS you choose determines the channel width (20 MHz vs. 40 MHz) and the number of spatial streams used by the mesh node. The default value is 1–23; the complete set of supported values. To specify a smaller range of values, enter a hyphen between the lower and upper values. To specify a series of different values, separate each value with a comma. Examples: 2–10 1,3,6,9,12 Range: 0–23.
VHT - Support MCS Map	A list of MCS values or ranges of values to be supported on this SSID. The MCS you choose determines the channel width (20 MHz vs. 40 MHz vs 80 MHz) and the number of spatial streams used by the mesh node. The default value is 1–23; the complete set of supported values. To specify a smaller range of values, enter a hyphen between the lower and upper values. To specify a series of different values, separate each value with a comma. Examples: 2–10 1,3,6,9,12 Range: 0–23.
Maximum VHT MPDU size	Maximum size of a VHT MPDU, in bytes. Range: 3895, 7991, 11454
Maximum number of MSDUs in an A-MSDU on best-effort AC	Maximum number of MSDUs in a TX A-MSDU on best-effort AC. TX-AMSDU disabled if 0. Range: 0-15 Default: 2
Maximum number of MSDUs in an A-MSDU on background AC	Maximum number of MSDUs in a TX A-MSDU on background. TX-AMSDU disabled if 0. Range: 0-15 Default: 2
Maximum number of MSDUs in an A-MSDU on video AC	Maximum number of MSDUs in a TX A-MSDU on video AC. TX-AMSDU disabled if 0. Range: 0-15 Default: 2
Maximum number of MSDUs in an A-MSDU on voice AC	Maximum number of MSDUs in a TX A-MSDU on voice AC. TX-AMSDU disabled if 0. Range: 0-15 Default: 0

6. Click **Submit**.

7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The profile name appears in the Mesh High-throughput SSID Profile list with your configured settings. To edit a mesh high-throughput SSID profile, select a profile from the Mesh High-throughput SSID Profile list and change the settings as desired.

Assigning a Profile to an AP Group

The following procedure describes how to assign a profile to an AP group:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > AP Groups**.
2. Select an AP group.
3. Click the **Profiles** tab.
4. Under the **Profiles** list, expand **Mesh**, and then select a **Mesh High-throughput SSID** profile.
5. In the **Mesh High-throughput SSID Profile** window, select a profile from the **Mesh High-throughput SSID profile** drop-down list.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

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.

Deleting a Profile

You can delete a mesh high-throughput SSID profile only if no APs or AP groups are associated with that profile.

The following procedure describes how to delete a mesh high-throughput SSID profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System >Profiles** tab.
2. Under the **All Profiles** menu, expand the **Mesh** tab and then select **Mesh High-throughput SSID**.
3. Select the Mesh High-throughput SSID profile you want to delete, and then click the **delete** icon.

Managing Mesh High-Throughput SSID Profiles in the CLI

You must be in config mode to create, modify or delete a mesh high-throughput SSID radio profile using the CLI. Specify an existing high-throughput SSID profile with the <profile-name> parameter to modify an existing profile, or enter a new name to create an entirely new profile.

This section contains the following topics:

- Creating or Modifying a Profile
- Assigning a Profile to an AP or AP Group
- Viewing High-throughput SSID Settings
- Deleting a Profile

Creating or Modifying a Profile

The following CLI commands create or modify a high-throughput SSID profile. If you do not specify a parameter for a new profile, that profile uses the default value for that parameter. Put the **no** option

before any parameter to remove the current value for that parameter and return it to its default setting. Enter **exit** to leave the high-throughput radio profile mode.

```
(host) [mynode] (config) #ap mesh-ht-ssid-profile <profile-name>
40mhz-enable
80mhz-enable
ba-amsdu-enable
clone
high-throughput-enable
ldpc
legacy-stations
max-rx-a-mpdu-size
max-tx-a-mpdu-size
max-tx-a-msdu-count-be
max-tx-a-msdu-count-bk
max-tx-a-msdu-count-vi
max-tx-a-msdu-count-vo
max-vht-mpdu-size
min-mpdu-start-spacing
mpdu-agg
no
short-guard-intvl-20mhz
short-guard-intvl-40mhz
short-guard-intvl-80mhz
stbc-rx-streams
stbc-tx-streams
supported-mcs-set
temporal-diversity
vht-supported-mcs-map
```

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.

```
(host) [mynode] (config) #ap mesh-ht-ssid-profile <profile-name> clone <source-
profile-name>
```

Assigning a Profile to an AP or AP Group

The following CLI command associates a mesh high-throughput SSID profile with an AP group:

```
(host) [mynode] (config) #ap-group <group> mesh-ht-ssid-profile <profile-name>
```

The following CLI command associates a mesh radio profile with an individual AP:

```
(host) [mynode] (config) #ap-name <name> mesh-ht-ssid-profile <profile-name>
```

Viewing High-throughput SSID Settings

The following CLI command displays a complete list of high-throughput profiles and their status:

```
(host) [mynode] (config) #show ap mesh-ht-ssid-profile
```

The following CLI command displays the settings of a specific high-throughput profile:

```
(host) [mynode] (config) #show ap mesh-ht-ssid-profile <profile-name>
```

Deleting a Profile

If no AP or AP group is using a mesh high-throughput SSID profile, you can delete that profile using the **no** parameter:

```
(host) [mynode] (config) #no ap mesh-ht-ssid-profile <profile-name>
```

Configuring Mesh Cluster Profiles

The mesh cluster configuration gets pushed from the controller to the mesh portal and the other mesh points, which allows them to inherit the characteristics of the mesh cluster of which they are a member. Mesh nodes are grouped according to a mesh cluster profile that contains the MSSID, authentication methods, security credentials, and cluster priority. Cluster profiles (including the default cluster profile) are not applied until you provision your APs for mesh.

This section contains the following topics:

- Managing Mesh Cluster Profiles in the WebUI
- Managing Mesh Cluster Profiles in the CLI

Managing Mesh Cluster Profiles in the WebUI

Use the following procedures to manage your mesh cluster profiles.

This section contains the following topics:

- Creating a Profile
- Associating a Mesh Cluster Profile to Mesh APs
- Editing a Mesh Cluster Profile
- Deleting a Mesh Cluster Profile

Creating a Profile

The following procedure describes how to create a mesh cluster profile:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Under the **All Profiles** pane, expand **Mesh** and click **Mesh Cluster profile**.
3. Click + in **Mesh cluster profile: New profile.**
4. Enter a **Profile name**.
5. Configure the mesh cluster settings described in Table 140.

Table 140: *Mesh Cluster Profile Configuration Parameters*

Parameter	Description
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 Aruba-mesh. Use the Cluster Name parameter to define a new, unique MSSID before you assign APs or AP groups to the mesh cluster profile. If you want a mesh cluster to use WPA2-PSK-AES encryption, do not use spaces in the mesh cluster name, as this may cause errors in mesh points associated with that mesh cluster. To view existing mesh cluster profiles, use the CLI command **show ap mesh-cluster-profile**. A mesh portal chooses the best cluster profile and provisions it for use. A mesh point can have a maximum of 16 cluster profiles. Default: Mesh cluster named Aruba-mesh.
RF Band	Indicates the band for mesh operation for multi-band radios. Select **a** or **g**. Important: If you create more than one mesh cluster profile for an AP or AP group, each mesh cluster profile must use the same band.

Parameter	Description
WPA Hexkey	Configures a WPA-PSK. This key must be of 64 hexadecimal characters. Re-enter the key in the Retype text box.
WPA Passphrase	Sets the WPA password that generates the PSK. The passphrase must be between 8–63 characters, inclusive. Re-enter the password in the Retype text box.
Encryption	Configures the data encryption, which can be either opensystem (no authentication or encryption) or wpa2-psk-aes. It is recommended to select wpa2-psk-aes and using the WPA-passphrase parameter to select a passphrase. Default: opensystem.

6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Associating a Mesh Cluster Profile to Mesh APs

The following procedure describes how to associate a mesh cluster profile to a group of mesh APs or an individual mesh AP. If you configure multiple cluster profiles with different cluster priorities, you manually override the link metric algorithm because the priority takes precedence over the path cost. In this scenario, the mesh portal uses the profile with the highest priority to bring-up the mesh network.

1. In the **Managed Network** node hierarchy, navigate to **Configuration > AP Groups**.
2. Select an AP group.
3. Navigate to the **AP group <Group Name> Mesh** tab and expand the **Mesh Cluster** profile.
4. To add a Mesh Cluster profile, click **+ in** the Mesh cluster window,
   a. To add an existing cluster, click **Add existing cluster** and select a cluster from the cluster name table. Click **Submit**.
   b. To add a new cluster, click **Create new cluster** and enter a **Cluster name**. Click **Submit**.

Editing a Mesh Cluster Profile

The following procedure describes how to edit a mesh cluster profile. If you modify any mesh cluster profile setting, you must reprovision your AP. For example, if you change the priority of a cluster profile from 5 to 2, you must reprovision the AP before you can assign priority 5 to another cluster profile. Reprovisioning the AP causes it to automatically reboot. For more information, see Provisioning Mesh Nodes on page 774.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Under the **All Profiles** pane, expand **Mesh > Mesh Cluster profile** and then select the profile you want to edit.
3. Change the mesh cluster settings as desired. Table 140 describes the parameters you can configure for new and existing mesh clusters.



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 **Save**, then remove the encryption type you no longer want and click **Save** again. You cannot delete one encryption type and add a different type in a single step.

4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy 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.

The following procedure describes how to delete a mesh cluster profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles** tab.
2. Under the **All Profiles** pane, expand **Mesh > Mesh Cluster profile**.
3. Click the **Delete** icon next to the name of the profile you want to delete.

Managing Mesh Cluster Profiles in the CLI

You must be in config mode to create, modify or delete a mesh cluster profile using the CLI. Specify an existing mesh cluster profile with the <profile-name> parameter to modify an existing profile, or enter a new name to create an entirely new profile. If you do not specify a parameter for a new profile, that profile uses the default value for that parameter.

Use the **no** option before any parameter to remove the current value for that parameter and return it to its default setting. Enter **exit** to leave the mesh cluster profile mode.

```
(host) [mynode] (config) #ap mesh-cluster-profile <profile>
clone <profile>
cluster <name>
no ...
opmode [opensystem | wpa2-psk-aes]
rf-band {a | g}
wpa-hexkey <wpa-hexkey>
wpa-passphrase <wpa-passphrase>
```

Creating Mesh Cluster Profile

The following CLI commands create and configure the mesh cluster profiles **cluster1** and **cluster2**.

```
(host) [mynode] (config) #ap mesh-cluster-profile cluster1
cluster corporate
opmode wpa2-psk-aes
wpa-passphrase mesh_123
rf-band a
(host) [mynode] (config) #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.

```
(host) [mynode] (config) #ap mesh-cluster-profile <profile-name> clone
<source-profile-name>
```

The following sections provide information on the various mesh cluster profile settings:

Viewing Mesh Cluster Profile Settings

The following command displays a complete list of mesh cluster profiles and their status:

```
(host) [mynode] (config) #show ap mesh-cluster-profile
```

The following command displays the settings of a specific mesh cluster profile:

```
(host) [mynode] (config) #show ap mesh-cluster-profile <profile-name>
```

Associating Mesh Cluster Profiles

For deployments with multiple mesh clusters, you must also configure 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.

The following CLI command associates a mesh cluster profile to an AP group in a single-cluster deployment:

```
(host) [mynode] (config) #ap-group <group> mesh-cluster-profile <profile-
name>
```

The following CLI command associates a mesh cluster profile to an individual AP in a single-cluster deployment:

```
(host) [mynode] (config) #ap-name <name> mesh-cluster-profile <profile-name>
```

The following CLI command associates a mesh cluster profile to an AP group in a multiple-cluster deployment:

```
(host) [mynode] (config) #ap-group <group> mesh-cluster-profile <profile-
name> priority <priority>
```

The following CLI command associates a mesh cluster profile to an individual AP in a multiple-cluster deployment:

```
(host) [mynode] (config) #ap-name <name>
mesh-cluster-profile <profile-name> priority <priority>
```

Example:

```
(host) [mynode] (config) #ap-group group1
mesh-cluster-profile cluster1 priority 5
mesh-cluster-profile cluster2 priority 10
(host) [mynode] (config) #ap-group2
mesh-cluster-profile cluster1 priority 10
mesh-cluster-profile cluster2 priority 5
mesh-radio-profile channel2
```

Excluding a Mesh Cluster Profile from a Mesh Node

The following CLI command excludes a specific mesh cluster profile from an AP:

```
(host) [mynode] (config) #ap-name <name> exclude-mesh-cluster-profile-ap
<profile-name>
```

Deleting a Mesh Cluster Profile

If no APs are using a mesh cluster profile, you can delete that profile using the **no** parameter in the CLI:

```
(host) [mynode] (config) #no ap mesh-cluster-profile <profile-name>
```

Configuring Mesh Clusters Associated with AP Groups

Mesh clusters are similar to an 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.

This section contains the following topics:

- Creating Mesh Cluster in the WebUI
- Editing Mesh Cluster in the WebUI
- Disassociating Mesh Cluster from AP Group in the WebUI
- Configuring Mesh Radio Settings in the WebUI
- Configuring Mesh High Throughput Associated with AP Groups

Creating Mesh Cluster in the WebUI

The following procedure describes how to create a mesh cluster using the WebUI:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > AP Groups**.
2. Select an AP Group from the **AP Groups** table.
3. In the **AP groups <Group name>** window, click the **Mesh** tab.
4. Click + in the **Mesh Clusters** table. The **Add Mesh Cluster** window opens.
   a. To add an existing cluster, select **Add Existing Cluster** and select an existing cluster.
   b. To create a new cluster, select **Create new cluster**, and then enter a cluster name.
5. Click **Submit**.



Optionally, drag and drop a mesh cluster to change its priority.

6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Editing Mesh Cluster in the WebUI

The following procedure describes how to edit a mesh cluster:

Not all mesh clusters can be edited. Only a mesh cluster with the least priority can be edited.



A mesh cluster with the least priority can be identified by a green-colored circle in the **Priority** column of the **Mesh Clusters** table.

To change the priority of a mesh cluster, drag and drop a row in **Mesh Clusters** table.

1. In the **Managed Network** node hierarchy, navigate to **Configuration > AP Groups**.
2. Select an AP Group in the **AP Groups** table and click **Mesh** tab in the **AP Groups<Group name>** window.
3. Select a mesh cluster in the **Mesh Clusters** table.
4. In the Mesh Cluster table:
   a. Enter a **Cluster name**.
   b. Select an encryption type.
   c. Select a band.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Disassociating Mesh Cluster from AP Group in the WebUI

The following procedure describes how to disassociate a mesh cluster from an AP group:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > AP Groups**.
2. Select an AP Group in the **AP Groups** table and click **Mesh** tab in the **AP Groups<Group name>** window.
3. Select a mesh cluster in the **Mesh Clusters** table.
4. Click the **Delete** icon next to the selected mesh cluster.
5. In the **Remove Mesh Cluster** window, click **Remove**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Configuring Mesh Radio Settings in the WebUI

The mesh radio settings determine the settings used by mesh nodes to establish mesh links and the path to the mesh portal.

The following procedure describes how to configure mesh radio settings associated with AP groups:

1. In the Managed Device node hierarchy, navigate to **Configuration > AP Groups**.
2. Select an AP Group from the **AP Groups** table.
3. In **AP groups<Group name>**, Click the **Mesh** tab.
4. Expand **Radio Settings**.
5. Configure the mesh high throughput parameters described in Table 141.

Table 141: *Mesh Radio Settings Configuration Parameters*

Parameter	Description
Minimum link SNR	The minimum link signal-to-noise ratio. The allowed range of values is between 1 and 255. The default value is 12.
Matrix algorithm	The algorithm used by a mesh node to select its parent. Use this parameter to optimize operation of the link matrix algorithm. Available options are: best-link-rssi: Select the parent with the strongest RSSI, regardless of the number of children a potential parent has. distributed-tree-rssi: Select the parent based on link-RSSI and node cost based on the number of children. This option evenly distributes the mesh points over high quality uplinks. Low quality uplinks are selected as a last resort. Default is distributed-tree-rssi.
Reselection mode	The reselection mode. Available options are: reselect-anytime reselect-never startup-subthreshold subthreshold-only
Heartbeat threshold	The maximum number of heartbeat messages that can be lost between neighboring mesh nodes. Range is 1–255. Default is 10.

Parameter	Description
Link threshold	The minimum RSSI threshold below which a link is considered as 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). Range is 1 to 255. Default is 12.
Max children	The maximum number of children a mesh node can accept. Range is 1 to 64. Default is 64.
Max hop count	The maximum hop count from the mesh portal. Range is 1 to 32. Default is 8.
Mesh private VLAN	The VLAN ID for control traffic between a remote mesh portal and mesh nodes. Do not use this VLAN ID for user traffic. Range is 0 to 4094. Default is 0 (disabled).
Mesh survivability	This feature is currently not supported and should only be enabled under the supervision of Aruba support.
Retry limit	The number of times a mesh node can re-send a packet. Range is 0 to 15. Default is 8.
RTS threshold	The packet size sent by mesh nodes. Mesh nodes transmitting frames larger than this threshold must issue RTS and wait for other mesh nodes to respond with CTS to begin transmission. This helps prevent mid-air collisions. Range is 256 to 2346. Default is 2333.

6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Configuring Mesh High Throughput Associated with AP Groups

The mesh high throughput settings define settings unique to 802.11n and 802.11ac-capable, high throughput APs.

The following procedure describes how to configure mesh high throughput associated with AP groups:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > AP Groups**.
2. Select an AP Group from the **AP Groups** table.
3. In the **AP Groups <Group Name>**, click the **Mesh** tab.
4. Expand **High Throughput**.
5. Configure the mesh high throughput parameters described in Table 142.

Table 142: *Mesh High Throughput Configuration Parameters*

Parameter	Description
High throughput (802.11n)	Toggle switch to enable or disable high throughput for 802.11n capable access points. Default is enabled.
Use 40 MHz channels	Enable or Disable the use of 40 MHz channels. Default is enabled. This parameter is not available if the **High Throughput(802.11n)** parameter is disabled.

Parameter	Description
Explicit transmit beamforming	Enable or Disable the use of explicit transmit beamforming. Default is enabled. This option is only available if **High throughput(802.11n)** is enabled.
Very high throughput (802.11ac)	Toggle switch to enable or disable high throughput for 802.11ac capable access points. Default is enabled.
Use 80 MHz channels	Enable or Disable the use of 80 MHz channels. Default is enabled. This option is only available if the **Very high throughput(802.11ac)** parameter is enabled.

6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Configuring Ethernet Ports for Mesh

If you use mesh to join multiple Ethernet LANs, configure and enable bridging on the mesh point Ethernet port This section describes how to configure Ethernet ports for bridging or secure jack operation using the wired AP profile. The wired AP profile controls the configuration of the Ethernet port (s) on your AP.



Mesh nodes only support bridge mode and tunnel mode on their wired ports (Ethernet 0 or Ethernet 1). Split tunnel mode is not supported. Use bridge mode to configure bridging on the mesh point Ethernet port. Use tunnel mode to configure secure jack operation on the mesh node Ethernet port.

When configuring the Ethernet ports on dual-port APs, note the following requirements for the AP configures as a mesh portal:

- Connect Ethernet 0 to the managed device to obtain an IP address. The wired AP profile controls Ethernet 1.
- Only Ethernet 1 supports secure jack operation.

This section contains the following topics:

- Configuring Bridging on the Ethernet Port
- Configuring Ethernet Ports for Secure Jack Operation
- Extending the Life of a Mesh Network

Configuring Bridging on the Ethernet Port

The following procedure describes how to configure bridging on the Ethernet port:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Under the **All Profiles** pane, expand **AP**.
3. Expand **Wired AP** and select a profile.

4. Under **General** settings, do the following:
   a. Enable the **Wired AP** 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 managed device and they appear on the configured access mode VLAN. Tagged packets are dropped. All packets received from the managed device and sent through 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 managed device. Untagged packets are forwarded to the managed device on the configured Native VLAN. Packets received from the managed device 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 **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands configure bridging on the Ethernet port:

```
 (host) [mynode] (config) #ap wired-ap-profile <profile>
broadcast
clone
forward-mode {bridge | split-tunnel | tunnel}
wired-ap-enable
```

Optionally, the following wired AP profile settings can be configured:

```
(host) [mynode] (config) #ap wired-ap-profile <profile>
switchport mode {access | mode | trunk}
switchport access vlan <vlan>
switchport trunk native vlan <vlan>
switchport trunk allowed vlan <vlan>
trusted
```

Configuring Ethernet Ports for Secure Jack Operation

You can configure the Ethernet port(s) on mesh nodes to operate in tunnel mode. Known as secure jack operation for mesh, this configuration allows Ethernet frames coming into the specified wired interface to be GRE tunneled to the managed device. 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 managed device 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 Ethernet 0, to tunnel the frame to the managed device.

When configuring mesh Ethernet ports for secure jack operation, note the following guidelines:

- Mesh points support secure jack on Ethernet 0 and Ethernet 1.
- Mesh portals only support secure jack on Ethernet 1. This function is only applicable to Aruba APs that support a second Ethernet port and mesh.

You configure secure jack operation in the wired AP profile.



The parameters in the wired AP profile only apply to the wired AP interface to which they are assigned. Two wired interfaces can have different parameter values.

The following procedure describes how to configure ethernet ports for secure jack operation:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles** tab.
2. Under the **All Profiles** pane, expand **AP.**
3. Open **Wired AP** and select a profile. The settings for the currently selected wired AP profile appear.
4. Under **General** settings, do the following:
   a. Select the **Wired AP enable** check box. This option is not selected by default.
   b. From the **Wired AP mode** drop-down list, select **normal** or **daisy-chain**.
   c. From the **Forward mode** drop-down list, select **tunnel**.
   d. Optionally, select **Trusted** to configure this as a trusted port.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands configure ethernet ports for secure jack operation:

```
(host) [mynode] (config) #ap wired-ap-profile <profile>
forward-mode tunnel
wired-ap-enable
```

Optionally, following wired AP profile settings can be configured:

```
(host) [mynode](config) #ap wired-ap-profile <profile>
trusted
```

Extending the Life of a Mesh Network

To prevent your mesh network from going down in the event of a managed device failure, modify the following settings in the AP system profile(s) used by mesh nodes to maintain the mesh network until the managed device is available:



It is recommended to use the default maximum request retries and bootstrap threshold settings for most mesh networks; however, if you must keep your mesh network alive, you can modify the settings as described in this section. The modified settings are not applicable if mesh portals are directly connected to the managed device.

- Maximum request retries: maximum number of times to retry AP-generated requests. The default is 10 times. If you must modify this setting, it is recommended to set a value of 10000.

- Bootstrap threshold: number of consecutive missed heartbeats before the AP rebootstraps. (Heartbeats are sent once per second.) The default is 9 missed heartbeats. If you must modify this setting, it is recommended to set a value of 5000.

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

The following procedure describes how to modify the AP system profile:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Under the **All Profiles** pane, expand **AP> AP System** and select the AP profile you want to edit.
3. In the **AP system profile: Name**, expand the **Advanced** settings and make the following changes:
   - Change the **Maximum Request Retries** to 10000.
   - Change the **Bootstrap threshold** to 5000.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands modify the AP system profile:

```
(host) [mynode] (config) #ap system-profile <profile>
max-request-retries 10000
bootstrap-threshold 5000
```

Configuring a Mesh Access List

Starting from ArubaOS 8.7.0.0, mesh access list feature allows each AP to establish mesh links only with the whitelisted neighboring APs.

Mesh access list can be configured using either **Configuration > AP Groups** page or **Configuration > System > Profiles > Mesh** page. The following procedure describes how to configure a mesh access list using **Configuration > AP Groups** page:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > AP Groups**.
2. Select an AP Group. In the **AP Groups > (name of the AP Group)** table, select **Mesh**.
3. Expand the **Mesh Topology** accordion.
4. To create a new rule, click **+** in the **Specific Connection Rules** table.
5. The **Add Connection Rule** Table is displayed. Enter a Name for the rule. For **Apply to**, select an AP / any number of APs for which the rule has to be applied. Click **Next**.
6. For **Allow Mesh to**, select the APs that are allowed to be discovered by the AP/ APs you chose in the previous step.
7. Click **Finish**.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following procedure describes how to configure a mesh access list using **Configuration > System > Profiles > Mesh** page:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles > Mesh**.

2. Select **Mesh Accesslist** and to create a new profile, click **+**.
3. Enter a **Profile name.** Click **+** in the AP name table and select the APs.
4. For **Type**, select **allow** or **deny**.
5. Click **Finish**.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI commands configure a mesh access list:

```
(host) [mynode] (config) #ap mesh-accesslist-profile Guest
(host) [mynode] (Mesh Accesslist profile "Guest") #ap-name <name>
(host) [mynode] (Mesh Accesslist profile "Guest") #type allow
```

The following CLI commands associate a mesh access list to an AP group:

```
(host) [mynode] (config) #ap-group default
(host) [mynode] (AP group "default") #mesh-accesslist-profile Guest
```

Provisioning Mesh Nodes

Provisioning mesh nodes is similar to thin APs; however, there are some key differences. Thin APs establish a channel to the managed device 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 managed device. 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 managed device. To do this, you must first configure mesh cluster profiles for each mesh node prior to deployment. See Creating and Editing Mesh Radio Profiles on page 752 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. If you create and select multiple mesh cluster profiles, the AP is provisioned with those as well. If you have a dual-radio AP and configure one radio for mesh and the other as a thin AP, each radio is provisioned as configured.

Each radio provisioned in mesh mode can operate in one of two roles: mesh portal or mesh point. You explicitly configure the role, as described in this section. This allows the AP to know whether it uses the mesh link (over the mesh point or mesh portal) or an Ethernet link to establish a connection to the managed device.

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 Mobility Master. If the other mesh cluster profiles are unavailable, mesh nodes use the recovery profile to establish a link to the Mobility Master; data forwarding does not take place.



If you create a new mesh cluster profile for an existing deployment, you must re-provision the AP for the new profile to take effect. If you re-provision mesh nodes that are already operating, re-provision the most distant (highest hop count) mesh points first followed by the mesh portals. If you re-provision the mesh portal first, the mesh points may be unable to form a mesh link. Re-provisioning the AP causes it to automatically reboot. This may cause a disruption of service to the network.

Provisioning Caveats

Remember the following when provisioning APs for mesh:

- You must provision the AP before you install it as a mesh node in a mesh deployment. To provision the AP, it must be physically connected to the local network or directly connected to the Mobility Master. When connected and powered on, the AP must also be able to obtain an IP address from a DHCP server on the managed device or from the Mobility Master.
- Make sure that the provisioned mesh nodes form a connected mesh network before physically deploying the APs. For more information, see Verifying Your Mesh Network on page 776.
- In multi-node networks, save the mesh cluster configuration before provisioning the mesh nodes. To save the configuration in the WebUI, at the top of any window, click **Pending Changes > Deploy changes**. To save your configuration in the CLI, use the command **write memory**.

Radio Selection for Mesh Links

The radio used for the mesh link can be configured in dual -5 GHz or split-5 GHz enabled access points. When dual 5 GHz radio or split 5 GHz radio is enabled on the access point, the operations on the 5 GHz band is split and carried out by two separate radios — lower 5 GHz radio and upper 5 GHz radio. The lower 5 GHz radio operates on channels 32–64 and the upper 5 GHz radio operates on channels 100-173. With two active 5 GHz radios, the mesh link functions can be dedicated to one radio while the other radio can be used to service clients.



This feature is currently supported only on 340 Series and 550 Series access points.

The radio used for the mesh link can be configured using the **rf-split5G-band-range** command and can be configured only using the CLI. This configuration can only be applied on dual-5 GHz radio or split- 5 GHz radio enabled APs. Apply the configuration and reboot the AP for the changes to take effect.

The following CLI command configures the radio for mesh link:

```
(host) [mynode] (config) #ap mesh-cluster-profile cluster1
(host) [mynode] (Mesh Cluster profile "cluster1") rf-split5G-band-range { first | full
| lower | upper }
```

The radio assignment and operating band information is listed in the following table:

Table 143: *Radio Assignment and Band Information*

Radio Mode	Radio	Operating Band
Dual-5 GHz	Radio 0	Lower 5 Ghz band
	Radio 1	Upper 5 Ghz band

Radio Mode	Radio	Operating Band
Split-5 GHz	Radio 0	Upper 5 Ghz band
	Radio 2	Lower 5 Ghz band

Verifying Your Mesh Network

To view a list of mesh APs in the WebUI, navigate to the one of the following pages in the **Managed Network** node hierarchy:

- **Dashboard > Network**
- **Dashboard > Controllers**

The following CLI commands display the mesh APs and the mesh topology tree using the CLI, issue the following commands:

```
(host) [mynode] #show ap mesh active
(host) [mynode] #show ap mesh topology
```

Verification Checklist

After provisioning the mesh APs, follow the steps below to ensure that the mesh network is up and operating correctly.

- Issue the command **show ap mesh topology** to verify all the mesh APs are up and the topology is as expected. (Wait 10 minutes after startup for the topology to stabilize.)
- Verify each mesh node has the expected RSSI to its neighboring mesh nodes. The mesh topology is updated periodically, so access the command-line interface and issue the command **show ap mesh neighbors**for the current status. If the RSSI is low, verify that the tx-power settings in the mesh node's 802.11a or 802.11g radio profiles are correct, or, if ARM is used, verify the correct minimum tx-power setting.
- Issue the command **show ap mesh debug provisioned-clusters** to verify that the mesh clusters are correctly defined and provisioned (with encryption if desired). Issue the **show running-config | include recovery** command to verify that the cluster's recovery profile matches the managed device's recovery profile.
- Verify antenna provisioning by issuing the **show ap provisioning** command and verify installation parameters for non-default installations (that is, standard indoor APs deployed outside, or outdoor APs deployed inside). Ensure all APs use the same channel list by issuing the **show ap allowed-channels** command.
- If the mesh-radio is to be reserved exclusively for mesh backhaul traffic, issue the command **show ap profile-usage** to identify the radio's 802.11a or 802.11g radio profile, then issue the command **show rf dot11a-radio-profile <profile>** or **show rf dot11g-radio-profile <profile>** to verify the radio is disabled in the profile. Next, use the **show ap bss-table** command to that verify no access Virtual APs are up on the mesh radio.

Configuring Remote Mesh Portals

You can configure a Remote Mesh portal using the WebUI or CLI.

This section contains the following topics:

- Creating a Remote Mesh Portal In the WebUI
- Provisioning a Remote Mesh Portal in the CLI

Creating a Remote Mesh Portal In the WebUI

A remote mesh portal must be provisioned as both a remote access point and a mesh portal. For instructions on provisioning the remote mesh portal as a remote access point, see Configuring the Secure Remote Access Point Service.

Wired ports on remote mesh portals can be configured in either bridge or split-tunnel forwarding mode. However, there are limitations to the forwarding modes that can be used by other mesh node types. Do not use bridge or split-tunnel forwarding mode for wired ports on mesh points. Virtual APs on remote mesh portals and remote mesh points also do not support bridge or split-tunnel forwarding mode.



A remote mesh portal does not support bridge mode Virtual APs or offline Virtual APs.

The following procedure describes how to create a remote mesh portal:

Provisioning the AP

The following procedure describes how to provision an AP:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Access Points**.
2. Open the **Remote APs** tab.
3. Select the AP to provision as a remote mesh portal, and then click **Provision**.
4. In the **Authentication** section, select the **Remote AP** option.
5. 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 IKE PSK.
6. In the **Master Discovery** section, set the Master IP address as the controller IP address.
7. In the **IP settings** section, select **Obtain IP Address Using DHCP**.
8. In the **AP List** section at the bottom of the window, click the **Mesh Role** drop-down list and select **Remote Mesh Portal**.

Defining the Mesh Private VLAN in the Mesh Radio Profile

The following procedure describes how to choose a new, non-zero tag value for the mesh private VLAN. Make sure that the mesh private VLAN so that it does not conflict with any local tags assigned in the mesh network. Once configured, all mesh points come up in that Mesh Private VLAN. This mesh private VLAN must not be used as a VLAN for any other virtual AP.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Under the **All Profiles** pane, expand **Mesh** and select **Mesh Radio**.
3. To edit an existing mesh profile, select the profile that you want to edit from **Mesh > Mesh Radio**.
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 **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy 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.

Assigning the Mesh Radio Profile to a Remote Mesh AP

The following CLI commands associate a mesh radio profile with an individual AP:

```
(host) [mynode] (config) #ap-name <name>
mesh-radio-profile <profile-name> priority <priority>
```

Assigning an RF Management Profile to a Remote Mesh AP

The following procedure describes how to manage the most common RF management settings.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > AP Groups** tab.
2. Select the name of an AP group from the AP groups tables.
3. Click the **Radio** tab below the AP groups tables to display the AP groups radio settings. The radio settings are divided into three sections, **Basic, Advanced**, and **Client Control.**
4. Modify the desired settings, then click **Submit.**
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Assigning a Mesh Cluster Profile to a Remote Mesh AP

The following CLI command associates a mesh cluster profile to an individual AP in a single-cluster deployment:

```
(host) [mynode] (config) #ap-name <name> mesh-cluster-profile <profile-name>
```

The following CLI command associates a mesh cluster profile to an individual AP in a multiple-cluster deployment:

```
(host) [mynode] (config) #ap-name <name>
mesh-cluster-profile <profile-name> priority <priority>
```



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.

Configuring a DHCP Pool

In this next step, you must configure a DHCP pool where the DHCP server is on the subnet associated with mesh private VLAN. Mesh points get their IP address from this subnet pool.

Configuring the VLAN ID of the Virtual AP Profile

Follow the procedure described below to configure the VLAN ID of the remote mesh AP's SSID profile. The VLAN of this Virtual AP must have the same VLAN ID as the mesh private VLAN.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. From the **All Profiles** list, select **Wireless LAN> Virtual AP**.
3. To edit an existing Virtual AP profile, select the Virtual AP profile you want to edit. To create a new Virtual AP profile, click **+** and enter a name for the new Virtual AP profile in the **Profile name** field. The Virtual AP profile settings are divided into four sections, **General, RF, Advanced,** and **Broadcast/Multicast**.
4. Under **General**, enter the **VLAN ID**.

5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

Provisioning a Remote Mesh Portal in the CLI

Reprovisioning the AP causes it to automatically reboot. When you use the CLI to reprovision a mesh node, you may also provision other AP settings.

The following CLI commands provision a remote mesh portal :

```
(host) [mynode] (config) #provision-ap
read-bootinfo ap-name <name>
mesh-role remote-mesh-portal
reprovision ap-name <name>
```

Chapter 29
Increasing Network Uptime With Redundancy Services

Aruba offers a variety of solutions to provide inter-controller redundancy in a centralized architecture to ensure highly available networks. These solutions can be used by an administrator to implement a highly redundant network depending on the requirements and constraints for each network design. ArubaOS high availability and VRRP redundancy features allow network administrators to significantly reduce network downtime and client traffic disruption during network upgrades or unexpected failures.

This chapter describes the various redundancy solutions such Mobility Masterredundancy, AP High Availability and VRRP redundancy. The following topics describe the procedures to configure various redundancy services and features:

- Mobility Master Redundancy Methods on page 780
- AP and User Redundancy Methods on page 794

Mobility Master Redundancy Methods

Aruba supports the following methods to configure a redundant Mobility Master:

- In the CLI on page 783
- Configuring Secondary Mobility Master Using Layer-3 Redundancy on page 788
- Configuring Mobility Master in VPNC Topology on page 793

ArubaOS now provides IPv6 network infrastructure support to configure Mobility Masters in Layer-2 and Layer-3 redundancy, as well as establish communication between Mobility Masters and managed devices in a VPNC topology.

Configuring Secondary Mobility Master Using Layer 2 Redundancy

The Mobility Master in the Aruba user-centric network acts as a single point of configuration for global policies such as firewall policies, authentication parameters, and RF configuration to ease the configuration and maintenance of a wireless network. To maintain a highly redundant network, the administrator can use another Mobility Master to act as a hot secondary for the primary Mobility Master using VRRP.

The primary and secondary Mobility Masters establish IPsec tunnel to securely synchronize data between them. To establish IPsec tunnel, the Mobility Masters must obtain the IP address of the peer Mobility Master for using Layer-2 Master redundancy configuration.



If your deployment is using a MM-VA 50 SKU and the topology is configured for Layer-2 redundancy, note that the standby Mobility Master will be counted as one Mobility Controller against the capacity and license count.

The topic includes the following sections:

- Before you Begin on page 781
- Configuring VRRP for Mobility Master on page 781

- Configuring Master Redundancy on page 785
- Configuring Database Synchronization on page 787

Before you Begin

Before you begin configuring VRRP redundancy, obtain the following network information:

- **VLAN ID** for the primary and secondary Mobility Master on the same Layer-2 network.
- **Virtual IP address** to be used for the VRRP instance.



Ensure that the two Mobility Masters are connected on the same broadcast domain (or Layer-2 connected) for VRRP operation. The two Mobility Masters must run the same version of ArubaOS.

Configuring VRRP for Mobility Master

The managed devices can now establish IPsec tunnels with the Mobility Masters containing both VRRP IPv4 and VRRP IPv6 addresses. However, the applications under the managed device use the IPv4 address to communicate with the Mobility Master. The managed devices establish an IPv6 tunnel with the VRRP IPv6 address of the Mobility Masters and the IPv4 packets are routed over IPv6 tunnel to the Mobility Master.

The following procedure configures VRRP on the Mobility Master.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Redundancy > L2 Redundancy** tab.
2. Expand the **Virtual Router Table** accordion.
3. Under **Virtual Router Table**, click **+** to add a new virtual router.
4. The **New Virtual Router** fields are displayed.
5. Select the IP version from the **IP Version** drop-down list.
6. Select the VLAN on which you want to configure VRRP from the **VLAN** drop-down list.
7. Set **Admin State** to **UP**.
8. Configure other VRRP parameters as described in Table 144.
9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.
12. Repeat steps 1-10 to configure VRRP on the other Mobility Master in the primary and secondary redundant pair.



Ensure to reload the device whenever you modify the Master VRRP configuration under Master Redundancy to avoid any configuration errors.

Table 144: *VRRP Configuration Parameters*

Parameter	Description
ID	The ID uniquely identifies this VRRP instance. For ease in administration, you should configure this with the same value as the VLAN ID.
Description	This is an optional text description to describe the VRRP instance.

Table 144: *VRRP Configuration Parameters*

Parameter	Description
IP version	Select IPv4 \ IPv6 from the drop-down list box.
Authentication Password	This is an optional password of up to eight characters that can authenticate VRRP peers in their advertisements. If this is not configured, there is no authentication password.
Retype authentication password	Reconfirm the password, if configured.
IP address	Based on the selection made in the IP version field, either IP Address \ IPv6 Address is displayed. This is the virtual IP address that will be owned by the elected VRRP master. Ensure that the same IP address and VRRP ID is used on each member of the redundant pair. **NOTE:** The IP address must be unique and cannot be the loopback address of the device. A maximum of only two virtual IPv6 addresses can be configured on each VRRP instance. Only IPv6 address format is supported for the v6 instance.
IPv6 address	Configure the virtual IPv6 address that will be owned by the elected VRRP master. Use the same IPv6 address on each member of the redundant pair. This IPv6 address will be redundant - it will be active on the VRRP master, and will become active on the VRRP backup when the VRRP master fails. **NOTE:** The IPv6 address must be unique and cannot be the loopback address of the device. A maximum of only two virtual IPv6 addresses can be configured on each VRRP instance. Starting from ArubaOS 8.2.1.0, you can configure a unique local address as the VRRP IPv6 address on the Mobility Master and the managed devices.
Priority	Priority level of the VRRP instance for the device. This value is used in the election mechanism for the master. When configuring VRRP on a standby device, use the default priority value of 100. For a master device, use a higher priority value, such as 110.
Advertisement interval (secs)	This is the interval, in seconds, between successive VRRP advertisements sent by the current master. The default interval time is recommended. Default: 1 second
Enable router Pre-emption	Selecting this option means that a device can take over the role of master if it detects a lower priority device currently acting as master.
Pre-emption delay (secs)	Specifying a value enables the delay timer. The timer is triggered when the VRRP state moves out of backup or init state to become a master. This is applicable only if you enable router pre-emption. When the timer is triggered, it forces VRRP to wait for a specified period of time, so that all the applications are ready before coming up. This prevents the APs from connecting to the managed device or the stand-alone controller before it can receive them. In the meantime, if there is an advertisement from another VRRP, the VRRP stops the timer and does not transition to master.
Hold Time	Hold time is the number of seconds until which a managed device waits, before starting VRRP to account for System or Network convergence delays. Default: 45 seconds

Table 144: *VRRP Configuration Parameters*

Parameter	Description
	NOTE: Configuring the hold time will not take affect if preemption is enabled.
Admin state	Administrative state of the VRRP instance. To start the VRRP instance, change the admin state to **UP** in the WebUI.
VLAN	VLAN on which the VRRP protocol runs.
Tracking master up-time	(Optional) Perform VRRP priority tracking based on how long the device has been the master. This feature is designed to ensure that a master will only be allowed to take and maintain control of the VRRP if it has been up for a certain amount of minutes (0-1440). This prevents an issue where a device that is periodically going up and down assumes the role of primary master.
Tracking master up-time priority	(Optional) The additional priority given to the master once it has been up for the time interval defined by the **Tracking Master Up-time** parameter.
Tracking VRRP master state ID	(Optional) Perform tracking based on the UP or DOWN state of another VRRP master by specifying the VRRP ID of the master to be tracked.
Tracking VRRP master state priority	(Optional) The priority taken away from a VRRP master if it is in a DOWN state. The priority levels are returned to their previous state when the VRRP master comes back up.
Tracking VLAN	(Optional) Perform VRRP priority tracking based on the UP or DOWN state of a VLAN. Click **+** below the **Tracking VLAN** table and specify the following values: ■ VLAN Id: ID of the VLAN to be tracked. ■ Subtract: Priority level to be subtracted from the device's VRRP priority if the tracked VLAN goes down.
Tracking interface	(Optional) Perform VRRP priority tracking based on the UP or DOWN state of a specific interface. Click **+** below the **Tracking Interface** table and specify the following values: ■ Interface: Interface Port to be tracked. ■ Subtract: Priority level to be subtracted from the device's VRRP priority if the tracked interface goes down.



You must configure both VRRP IPv4 and VRRP IPv6 addresses with the same priority in Mobility Masters so that either of the VRRP IP addresses can become active or standby during a failover. This is to ensure that a single Mobility Master acts as a master at any given point of time. Also, ensure that you shut down both VRRP IPv4 and VRRP IPv6 addresses in the primary Mobility Master in Layer-2 or Layer-3 redundancy, so that the APs are up and running after a failover.

Run the following CLI commands to configure VRRP IPv4 address on both the Mobility Masters:

```
(MM-Primary) [mynode] (config) #vrrp <id>
(MM-Primary) ^[mynode] (config-submode)#ip address <ip addr>
(MM-Primary) ^[mynode] (config-submode)#vlan <id>
(MM-Primary) ^[mynode] (config-submode)#description <string>
(MM-Primary) ^[mynode] (config-submode)#preempt delay <seconds>
(MM-Primary) ^[mynode] (config-submode)#priority <level>
(MM-Primary) ^[mynode] (config-submode)#no shutdown
```

Run the following CLI commands to configure VRRP IPv6 address on both the Mobility Masters:

```
(MM-Primary) [mynode] (config) #vrrp ipv6 <id>
(MM-Primary) ^[mynode] (config-submode)#ipv6 address <X:X:X:X::X>
(MM-Primary) ^[mynode] (config-submode)#vlan <id>
(MM-Primary) ^[mynode] (config-submode)#description <string>
(MM-Primary) ^[mynode] (config-submode)#preempt delay <seconds>
(MM-Primary) ^[mynode] (config-submode)#priority <level>
(MM-Primary) ^[mynode] (config-submode)#no shutdown
```

The following sample CLI commands configure virtual router 10 for IPv4 address on the initially-preferred Mobility Master:

```
(MM-Primary) [mynode] (config) #vrrp 10
(MM-Primary) ^[mynode] (config-submode)#ip address 192.168.10.245
(MM-Primary) ^[mynode] (config-submode)#vlan 1
(MM-Primary) ^[mynode] (config-submode)#description "Preferred-Master"
(MM-Primary) ^[mynode] (config-submode)#preempt delay 4
(MM-Primary) ^[mynode] (config-submode)#priority 200
(MM-Primary) ^[mynode] (config-submode)#no shutdown
```

The following sample is the corresponding VRRP configuration for IPv4 address for the backup Mobility Master:

```
(MM-Backup) [mynode] (config) #vrrp 10
(MM-Backup) ^[mynode] (config-submode)#ip address 192.168.10.245
(MM-Backup) ^[mynode] (config-submode)#vlan 1
(MM-Backup) ^[mynode] (config-submode)#description "Backup-Master"
(MM-Primary) ^[mynode] (config-submode)#preempt delay 4
(MM-Backup) ^[mynode] (config-submode)#priority 100
(MM-Backup) ^[mynode] (config-submode)#no shutdown
```

The following sample CLI commands configure virtual router 167 for IPv6 address on the initially-preferred: Mobility Master:

```
(MM-Primary) [mynode] (config) #vrrp ipv6 167
(MM-Primary) ^[mynode] (config-submode)#ipv6 address 2021:1:1:167::254
(MM-Primary) ^[mynode] (config-submode)#vlan 167
(MM-Primary) ^[mynode] (config-submode)#description "Preferred-Master"
(MM-Primary) ^[mynode] (config-submode)#preempt delay 1
(MM-Primary) ^[mynode] (config-submode)#priority 120
(MM-Primary) ^[mynode] (config-submode)#no shutdown
```

The following sample is the corresponding VRRP configuration for the backup: Mobility Master:

```
(MM-Backup) [mynode] (config) #vrrp ipv6 167
(MM-Backup) ^[mynode] (config-submode)#ipv6 address 2021:1:1:167::254
(MM-Backup) ^[mynode] (config-submode)#vlan 167
(MM-Backup) ^[mynode] (config-submode)#description "Backup-Master"
(MM-Primary) ^[mynode] (config-submode)#preempt delay 1
(MM-Backup) ^[mynode] (config-submode)#priority 90
(MM-Backup) ^[mynode] (config-submode)#no shutdown
```

Run the following CLI command on the Mobility Master (both primary and secondary) to verify the VRRP configuration for IPv4 address:

```
(MM-Primary) [mynode] (config) #show vrrp <vrid>
```

The following output is displayed on the primary Mobility Master for IPv4 address:

```
(MM-Primary) [mynode] (config) #show vrrp 10
Virtual Router 10:
Description
Admin State UP, VR State MASTER
IP Address 192.168.10.245, MAC Address 00:00:5e:00:01:34, vlan 1
Priority 200, Advertisement 1 sec, Preemption Enable Delay 0
Auth type NONE ********
tracking is not enabled
```

The following output is displayed on the backup Mobility Master for IPv4 address:

```
(MM-Backup) [mynode] (config) #show vrrp 10
Virtual Router 10:
Description
Admin State UP, VR State BACKUP
IP Address 192.168.10.245, MAC Address 00:00:5e:00:01:34, vlan 1
Priority 100, Advertisement 1 sec, Preemption Disable Delay 0
Auth type NONE ********
tracking is not enabled
```

Run the following CLI command on the Mobility Master (both primary and secondary) to verify the VRRP configuration for IPv6 address:

```
(MM-Primary) [mynode] (config) #show vrrp ipv6 <vrid>
```

The following output is displayed on the primary Mobility Master for IPv6 address:

```
(MM-Primary) [mynode] (config) #show vrrp ipv6
    Virtual Router 167:
    Description
    Admin State UP, VR State MASTER
    IPv6 Address 2021:1:1:167::254
    MAC Address 00:00:5e:00:02:a7, vlan 167
    Priority 120, Advertisement 1 sec, Preemption Enable Delay 1
    tracking is not enabled
```

The following output is displayed on the backup Mobility Master for IPv6 address:

```
(MM-Backup) [mynode] (config) #show vrrp ipv6
    Virtual Router 167:
    Description
    Admin State UP, VR State BACKUP
    IPv6 Address 2021:1:1:167::254
    MAC Address 00:00:5e:00:02:a7, vlan 167
    Priority 90, Advertisement 1 sec, Preemption Enable Delay 1
    tracking is not enabled
```

Use the following command to configure the logging level for VRRP configured on the Mobility Master:

```
(host) [mymode] (config)# logging system process vrrp level <category> [subcat
<subcategory>]
```

Use the following command to debug issues related to VRRP in the GSM channel:

```
(host) [mymode]# show gsm debug channel vrrp_info
```

Configuring Master Redundancy

ArubaOS now allows you to also configure IPv6 address of the peer Mobility Master to establish an IPv6 IPsec tunnel between the primary and secondary Mobility Masters in Layer-2 Master redundancy configuration. All the applications that are involved in data synchronization between the two Mobility Masters use the IPv6 address of the peer Mobility Master.

You can configure the Master redundancy either using WebUI or CLI:

1. In the active **Mobility Master** node hierarchy, navigate to the **Configuration > Redundancy > L2 Redundancy** tab.
2. Expand the **Master Redundancy** accordion.
3. Select **IPv4** or **IPv6** radio button in the **IP version** field based on your preference.
4. Enter the virtual router ID of the VRRP instance in the **Master VRRP** field.
5. In the **IP address of peer** field, enter the IP address of the peer Mobility Master for master redundancy.
6. In the **Authentication** field, select **IPSec Key** or **Certificate** from the drop-down list based on your preference.
7. (Optional) Select **IPSec Key** from the **Authentication** field drop-down list, and do the following:
   a. In the **IPSec key of peer** field, specify the IPsec authentication password.
   b. In the **Retype IPSec key** field, retype the IPsec authentication password entered in step 7(a).
8. (Optional) Select **Certificate** from the **Authentication** field drop-down list, and do the following:
   a. (Optional) In the **Certificate type** field, select **Custom** from the drop-down list,
      1. In the **Peer's MAC address** field, enter the peer MAC address of the certificate on the redundant Mobility Master.
      2. In the **CA certificate** field, enter the user-defined name of a trusted CA certificate installed on the redundant Mobility Master.
      3. In the **Server certificate** field, enter the user-defined name of a server certificate installed on the redundant Mobility Master.
      4. In the **Suite B algorithm** field, select **GCM 128** or **GCM 256** based on your preference.
   b. (Optional) In the **Certificate type** field, select **Factory** from the drop-down list.
      1. In the **Peer's MAC address** field, enter the MAC address of the certificate on the redundant Mobility Master.



Use the **show tpm cert-info** command to obtain the MAC address of TPM and factory-installed certificate, and the **show inventory** command to obtain the Management Port Hardware MAC address of PSK and custom certificate. The factory-installed certificate is applicable to hardware devices only.

9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.
12. Repeat steps 1-11 to configure the other Mobility Master.

Run the following commands on Mobility Master (both primary and secondary) to associate the VRRP instance for Layer-2 Master redundancy:

```
(MM-Primary) [mynode] (config) #master-redundancy
(MM-Primary) [mynode] (config-submode)#master-vrrp {<id>|ipv6 <id>}
(MM-Primary) ^[mynode] (config-submode)#peer-ip-address {<ip-addr>|ipv6} [ipsec
<KEY>|ipsec-custom-cert|ipsec-facttory-cert]
```

```
(MM-Primary) [mynode] (config) #write memory
The following sample commands configure Master redundancy on the primary Mobility Master
using IPv6 address:
(MM-Primary) [mynode] (config) #master-redundancy
(MM-Primary) [mynode] (config-submode)#master-vrrp ipv6 10
(MM-Primary) ^[mynode] (config-submode)#peer-ip-address ipv6 2001:1:2:2020::1 ipsec
aruba123
(MM-Primary) ^[mynode] (config) #write memory
```

The following sample commands are used for Master redundancy configuration on the secondary Mobility Master using IPv6 address:

```
(MM-Backup) [mynode] (config) #master-redundancy
(MM-Backup) [mynode] (config-submode)#master-vrrp ipv6 10
(MM-Backup) ^[mynode] (config-submode)#peer-ip-address ipv6 2001:1:2:2020::3 ipsec
aruba123
(MM-Backup) ^[mynode] (config) #write memory
Run the following sample commands to obtain the MAC address of TPM and factory-
installed certificate:
(MM-Primary) [mynode] #show tpm cert-info

=====================================
TPM manufacturing factory certificate
=====================================
subject= /CN=TWC7K35009::20:4c:03:0e:d2:84
issuer= /DC=com/DC=arubanetworks/DC=ca/CN=DEVICE-CA1
serial=5027A6B60000010D2D76
notBefore=Jan 13 09:16:16 2017 GMT
notAfter=Sep 14 03:21:14 2032 GMT
    Signature Algorithm: sha1WithRSAEncryption
```

Execute the following CLI command on the Mobility Master (both primary and secondary) to verify the Master redundancy configuration:

```
(MM-Primary) [mynode] #show master-redundancy
```

The following output is displayed on the primary Mobility Master:

```
(MM-Primary) [mynode] #show master-redundancy
Master redundancy configuration:
VRRP Id 10 current state is MASTER
Peer's IP Address is 192.168.10.244
Peer's IPSEC Key is ********
```

The following output is displayed on the secondary Mobility Master:

```
(MM-Backup) [mynode] #show master-redundancy
Master redundancy configuration:
VRRP Id 10 current state is BACKUP
Peer's IP Address is 192.168.10.243
Peer's IPSEC Key is ********
```

Configuring Database Synchronization

In a redundant Mobility Master scenario, you can configure a redundant pair to synchronize their WMS and local user databases. You can either manually or automatically synchronize the databases.

When manually synchronizing the database, the active VRRP master synchronizes its database with the secondary Mobility Master. The command takes effect immediately.

When configuring automatic synchronization, you can set how often the two Mobility Masters synchronize their databases. To ensure successful synchronization of database events, you must set the recommended default period of 30 minutes.

You can configure the database synchronization using either WebUI or CLI.

1. Navigate to **Mobility Master > Configuration > Redundancy** and expand the **Master Redundancy** accordion.
2. Click the **Database synchronization** toggle switch to enable this setting.
3. Enter the frequency of synchronizing the databases in the **Sync period** field.
4. The range is 1-25200 minutes, and the default value is 30 minutes.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.
8. Repeat steps 1-6 for the secondary Mobility Master.

Use the following command to configure database synchronization and the scheduled interval for synchronizing the databases:

```
(MM-Primary) [mm] (config) #database synchronize period <minutes>
Use the following command to verify the database synchronization  on the Mobility Master:
(MM-Primary)[mm] (config)#show database synchronize
```

Configuring Secondary Mobility Master Using Layer-3 Redundancy

ArubaOS 8.2.0.0 introduces support for a redundant pair of Mobility Masters in a layer 3 network. This prevents a scenario where a Mobility Master acts as a single point of failure if the link to the Mobility Master goes down, or a co-located secondary Mobility Master VRRP controller pair fail due to a network failure or local natural disaster.

The Mobility Masters in Layer-3 redundancy are situated across different Layer-2 networks and take the role of primary and secondary Mobility Masters. Similar to Layer-2 redundancy, the primary and secondary Mobility Masters establish IPsec tunnel to securely synchronize data between them by using Layer-3 Master redundancy configuration.



It is mandatory to use VRRP IPv6 address instead of switch IPv6 address to establish an IPsec tunnel, when VRRP IPv6 address is configured in both primary and secondary Mobility Masters. This is to ensure that there are no database synchronization failures between primary and secondary Mobility Masters in a Layer-3 redundancy.

ArubaOS now allows you to also configure IPv6 address of the peer Mobility Master to establish an IPv6 IPsec tunnel between the primary and secondary Mobility Masters in Layer-3 Master redundancy configuration. All the applications that are involved in data synchronization between the two Mobility Masters use the IPv6 address of the peer Mobility Master.



For Layer-3 redundancy to work, ensure that ICMP protocol is not blocked on managed devices, VPN concentrators, and Mobility Masters in the user network.

Configuring Layer-3 Redundancy

The Layer-3 redundancy feature will support Active-Standby Model. The Layer-3 redundancy role is driven by user configuration at both the primary and secondary Mobility Master. Once the systems are set up for Layer-3 redundancy, the switchover event will take place when the primary Mobility Master goes down. The secondary Mobility Master will provide the Mobility Master functionality without any user intervention.

Managed devices will have the management tunnel with only one Mobility Master at any given time. The managed device will try to connect to the secondary Mobility Master if it looses connectivity with the primary Mobility Master. The secondary Mobility Master will accept the management tunnel connections from a managed device only if its tunnel with primary Mobility Master is down. This will ensure that the Layer-3 switchover event is processed only if the primary Mobility Master is down and not due to flaky connectivity between the managed device and primary Mobility Master.

However, since the applications under the managed device use the IPv4 address to communicate with the Mobility Master, the managed devices establish an IPv6 tunnel with the primary or secondary Mobility Masters and the IPv4 packets are routed over IPv6 tunnel to the Mobility Master.

Listed below are the salient features of Layer-3 Redundancy:

- Configuration and database events are synced automatically from the primary to secondary Mobility Master.
- Managed devices detect a failure in the primary Mobility Master and automatically switch to the secondary Mobility Master after 15 minutes.
- The switchover event in the managed device will have minimal service impact, if any.
- Support for centralized licensing, a single license for both primary and secondary Mobility Masters.
- Layer-2 and Layer-3 redundancy will work together.
- When the primary Mobility Master comes back up all managed devices will switch back to primary Mobility Master with minimal service impact, if any.

The following procedure configures Layer-3 Master redundancy on the primary Mobility Master:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Redundancy > L3 redundancy** tab.
2. Select the **Controller role**.
3. Enter the **Sync period** (in hours).

   The minimum value is 2 hours and the maximum value is 24 hours.

   For Layer-2 redundancy, the minimum value is 1 minute and the maximum value is 25200 minutes.

4. Enter the **IP address of the peer.**
5. Select the authentication method from the **Authentication** drop-box.
   a. If **IPSec key** is selected as an authentication method, enter the **IPSec key of the peer** and **Re-type the key**
   b. If **Certificate** is selected as an authentication method and **Factory** is selected as the **Certificate type,** enter the**Peer's MAC address**
   c. If **Certificate** is selected as an authentication method and **Custom** is selected as the certificate type, enter the **Peer's MAC address, CA certificate, Server Certificate** and select a **Suite B algorithm** from the drop-down list.
6. Click **Submit**.

7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following procedure configures the details of the primary Mobility Master on a managed device for Layer 3 redundancy:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Controllers** tab.
2. Enable the **l3 redudancy** toggle switch.
3. Select **Direct** or **Via VPN concentrator** for **Connection to master**.
4. Select the **IP address version of master**.
5. Enter the **IPv4 address of master.**
6. Enter the **FQDN of the controller.**
7. Choose the Source interface from the **Source interface** drop down list.
8. Select the authentication method from the **Authentication** drop-box.
    a. If **IPSec key** is selected as an authentication method, enter the **IPSec key of the peer** and **Re-type the key**.
    b. If **Certificate** is selected as an authentication method and **Factory** is selected as the **Certificate type**, enter the **Peer's MAC address**.
    c. If **Certificate** is selected as an authentication method and **Custom** is selected as the **Certificate type**, enter the **Peer's MAC address**, **CA certificate**, **Server Certificate** and select a **Suite B algorithm** from the drop-down list.
9. Enter the **MAC address of master**.
10. Enter the **MAC address of the redundant master**.
11. Enable the **This controller is acting as VPN concentrator** check-box if the controller is a VPN concentrator.
12. Click **Submit**.
13. Click **Pending Changes**.
14. In the **Pending Changes** window, select the check box and click **Deploy** changes.

The following procedure configures the details of the secondary Mobility Master on a managed device for Layer 3 redundancy:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Controllers** tab.
2. Enable the **l3 redudancy** toggle switch.
3. Enter details for the primary Mobility Master.
4. Expand **Secondary Mobility Master**.
5. Select **Direct** or **Via VPN concentrator** for **Connection to master**.
6. Select the **IP address version of master**.
7. Enter the **IPv4 address of master.**
8. Enter the **FQDN of the controller.**
9. Choose the Source interface from the **Source interface** drop down list.
10. Select the authentication method from the **Authentication** drop-box.
    a. If **IPSec key** is selected as an authentication method, enter the **IPSec key of the peer** and **Re-type the key**.

b. If **Certificate** is selected as an authentication method and **Factory** is selected as the **Certificate type**, enter the **Peer's MAC address**.

c. If **Certificate** is selected as an authentication method and **Custom** is selected as the **Certificate type**, enter the **Peer's MAC address**, **CA certificate**, **Server Certificate** and select a **Suite B algorithm** from the drop-down list.

11. Enter the **MAC address of master**.
12. Enter the **MAC address of the redundant master**.
13. Enable the **This controller is acting as VPN concentrator** check-box if the controller is a VPN concentrator.
14. Click **Submit**.
15. Click **Pending Changes**.
16. In the **Pending Changes** window, select the check box and click **Deploy** changes.

The following CLI commands configure Layer-3 Master redundancy on the primary Mobility Master:

```
(MM-Primary) [mynode] (config) #master-l3redundancy
(MM-Primary) ^[mynode] (config-submode)#l3-peer-ip-address {<ip addr>|<ipv6>} ipsec
<KEY>
(MM-Primary) [mynode] (config) #write memory
The following CLI commands configure Layer-3 Master redundancy on the primary Mobility
Master:
(MM-Primary) [mynode] (config) #master-l3redundancy
(MM-Primary) ^[mynode] (config-submode)#l3-peer-ip-address 2001:1001::201 ipsec
aruba123
(MM-Primary) ^[mynode] (config) #write memory
The following sample is a   Layer-3 Master redundancy configuration on the backup
Mobility Master:
(MM-Backup) [mynode] (config) #seondary masteripv6 2001:1001::203
(MM-Backup) ^[mynode] (config-submode)#peer-ip-address 192.168.10.243 ipsec aruba123
(MM-Backup) ^[mynode] (config) #write memory
```



Ensure that the managed device has the same connection type between primary and secondary Mobility Masters to establish a tunnel using PSK, custom, or factory-installed certificate.

Important

Configuration changes cannot be made on the secondary Mobility Master. In a scenario where the primary Mobility Master is down and configuration changes need to be made on the secondary Mobility Master the user must change the sync state of the secondary Mobility Master to primary.

To preserve these configuration changes, a Layer-3 synchronization between the new primary Mobility Master and the old primary Mobility Master should take place. For the synchronization to take place the sync state of the old primary Mobility Master should be changed from primary to secondary state.

When the L3 sync state of a Mobility Master is changed from primary to secondary, the Mobility Master reboots to ensure a proper cleanup of the Mobility Master before new configurations or data is pushed from the new primary Mobility Master.

Once the roles of Mobility Masters are reversed, the user should ensure the managed devices point to the correct primary Mobility Master and secondary Mobility Master by changing the respective master IP address addresses.



The change of master IP and secondary master IP address that takes place on the primary Mobility Master from the managed device node should be done in the same write memory cycle. If this procedure is not done in same write memory cycle, the managed devices may point to the same IP as their primary and secondary Mobility Masters. If this happens reconfiguring the correct secondary masterip when the managed devices are up will fix the issue.

The change of master IP and secondary master IP address that takes place on the primary Mobility Master from the managed device node should be done in the same write memory cycle. If this procedure is not done in same write memory cycle, the managed devices may point to the same IP as their primary and secondary Mobility Masters. If this happens reconfiguring the correct secondary masterip when the managed devices are up will fix the issue.

Database Synchronization between Layer-3 Peers

Data synchronization between the Layer-3 peers only happens for the data that is already synchronized between Layer-3 pairs.

Table 145: *Layer-3 Synchronization*

	Synchronized Across Layer-3 Peers
Configuration	Yes
Database	Yes
Certificates	Yes
Captive Portal	Yes
Dynamic State Information of Services	No
Monitoring Data	No
Data Distribution Service State	No

Activate Provisioning

The Activate provisioning rule is enhanced to include the following data when Layer-3 Redundancy level is configured. Separate titles for primary data center and secondary data center are displayed to differentiate information.

```
Primary Data Center Secondary Data Center
================== ====================
Primary Master Controller: Primary Master Controller:
Master Controller IP: Master Controller IP:
Secondary Master Controller: Secondary Master Controller:
Primary VPN Concentrator MAC: Primary VPN Concentrator MAC:
VPN Concentrator IP: VPN Concentrator IP:
Secondary VPN Concentrator MAC: Secondary VPN Concentrator MAC:
```



The Activate provisioning rule does not provide the managed device with provisioning information.

Health Check Manager

The Health Check Manager provides detailed information on the health of uplinks. The Health Check Manager periodically pings and reports if the devices at the other end of the uplinks are reachable. Each

managed device interfaces with the Health Check Manager that provides information on the state of uplinks in both the primary and the secondary Mobility Master.



The secondary Mobility Master allows the connection of managed devices only if it determines that the primary Mobility Master is down for 15 minutes.

Load Balancing

The Load Balancing feature, provides a way to load balance across multiple uplinks between the same end points or across Mobility Masters.

Support is provided only for one active management tunnel. There is no support for load balancing the management traffic.

Configuring Mobility Master in VPNC Topology

In a VPNC based topology, the managed devices do not directly establish IPsec tunnel with Mobility Masters. The managed devices establish IPsec tunnel with VPNC which in turn establishes IPsec tunnel with the Mobility Master. ArubaOS now provides support to establish IPv6 tunnel between managed devices to VPNC and between VPNC to Mobility Masters. The VPNC acts as passthrough for both IPv4 as well as IPv6 communication between Mobility Master and managed device.



You can establish an IPv4 IPsec tunnel between a managed device and VPNC while having IPv6 IPsec tunnel between VPNC and Mobility Master in a network deployment. However, you cannot have an IPv6 IPsec tunnel between managed device and VPNC while having IPv4 IPsec tunnel between VPNC and Mobility Master.

The following CLI commands establish IPsec tunnel between VPNC and the managed device by adding peer details of the managed device on the VPNC:

```
(host)[mynode](config) #change-config-node /md
(host)[md](config) #vpn-peer peer-mac 00:0b:86:9a:6b:37 cert-auth factory-cert
(host)[md](config) #vpn-peer peer-mac 00:0b:86:9a:6b:37 pre-share-key aruba123
```

The following CLI commands add the Mobility Master details on VPNC or managed device:

```
(host)[md](config) #change-config-node /mm
(host)[mm](config) #local-peer-mac 00:0b:86:9a:6b:37 ipsec aruba@123
```

The following CLI command configures primary Mobility Master in VPNC topology:

```
(MM-Primary) [md] (config) #masteripv6 <ipv6-address> vpn-ipv6 <ipv6-address> ipsec
<KEY>| ipsec-factory-cert <options> | ipsec-custom-cert <options>] [ peer-mac-1 <MM
mac> ] [ peer-mac-2 <MM mac 2> interface vlan <id> masteripv4 <ipv4 address>
(MM-Primary) [md] (config) #write memory
```

The following sample CLI commands configure the primary Mobility Master:

```
(MM-Primary) [md] (config) #masteripv6 2001:1001::201 vpn-ipv6 2001:1001::204 ipsec
aruba123 peer-id <vpn mac address> interface vlan <vlan ID> masteripv4 <IPv4 address>
(MM-Primary) [md] (config) #write memory
```

The following sample CLI commands configure the secondary Mobility Master:

```
(MM-Backup) [md] (config) #secondary masteripv6 2001:1001::203
(MM-Backup) [md] (config-submode)##ipsec aruba123 vpn-ipv6 2001:1001::205 same as
above
(MM-Backup) [md] (config) #write memory
```

AP and User Redundancy Methods

Aruba supports the following AP and user redundancy methods:

- Controller Clustering on page 374—Cluster is a combination of multiple managed devices working together to provide high availability to all the clients and ensure service continuity when a failover occurs. More information on Controller Clustering and load balancing are covered in the chapter, Controller Clustering on page 374
- Active/Active Deployment Model on page 796—The High Availability WLAN redundancy solution enables campus APs to seamlessly failover to the standby controllers upon losing connectivity with the active controller. This significantly reduces the AP downtime.
- VRRP Redundancy on page 802—VRRP redundancy enables APs to failover to a backup controller when the AP's master becomes unavailable.



IPv6 Remote APs are not supported for clusters and high availability scenarios as well as mesh points.

AP High Availability Overview

The following topics in this section provides an overview about High Availability and VRRP redundancy:

- Learn more about High-Availability and VRRP Redundancy on page 794
- Controller Role Types on page 795
- AP Communication with Controllers on page 795
- Redundancy and High Availability Requirements and Limitations on page 795
- High Availability Deployment Models on page 796

For information to help you plan your redundancy solution, refer to the following topics under this section:

- High Availability with Extended Capacity on page 797
- Client State Synchronization on page 799
- High Availability Inter-Controller Heartbeats on page 799
- Configuring High Availability on page 799
- Migrating from VRRP or Backup-LMS Redundancy on page 801

Learn more about High-Availability and VRRP Redundancy



This section is applicable only for a stand-alone controller or a managed device on the Mobility Master.

When you enable the High Availability WLAN redundancy solution, campus APs that lose contact with their active controller do not need to re-bootstrap when they failover to the standby controller, significantly reducing AP downtime. APs using the High Availability features regularly communicate with the standby controller so the controller has a light workload to process in the event of an AP failover. This results in very rapid failover times and a shorter client reconnect period. Therefore, High Availability is usually preferable to other redundancy solutions (like a backup-LMS) that can put a heavy load on the backup controller during failover, which results in slower failover performance.



High Availability supports failover for campus APs using tunnel, or decrypt-tunnel, or bridge forwarding modes. It does not support failover for remote APs.

AP Fast Failover on bridge forwarding mode virtual AP is supported only on 7200 Series controllers.

Controller Role Types

A controller using this feature can have one of three high availability roles: **active**, **standby**, or **dual**. An active controller serves APs, but cannot act as a failover standby controller for any AP except those that it serves as an active controller. A standby controller acts as a failover backup controller, but cannot be configured as the primary controller for any AP. A dual controller can support both roles, acting as the active controller for one set of APs, and a standby controller for another set of APs.



A controller is assigned the **dual** role if no other role is specified.

AP Communication with Controllers

The High Availability features work across Layer-3 networks, so there is no need for a direct Layer-2 connection between controllers in a high-availability group.

When the AP first connects to its active controller, the active controller provides the IP address of a standby controller, and the AP attempts to establish a tunnel to the standby controller. If an AP fails to connect to the first standby controller, the active controller will select a new standby controller for that AP, and the AP will attempt to connect to that standby controller.

An AP will failover to its backup controller if it fails to contact its active controller through regular heartbeats and keepalive messages, or if the user manually triggers a failover using the WebUI or CLI. If inter-controller heartbeat is enabled, APs can failover even when the standby controller misses its heartbeats with the active controller.

Starting from ArubaOS 8.4.0.0, when an AP is disconnected without properly notifying the controller, the controller ages out the AP by detecting the AP heartbeats failure. Since the AP heartbeats are sent in clear text format, attackers can easily implement fake heartbeats and remove the valid APs without losing the AP connection state on the controller. To prevent attackers from sending rogue APs as valid APs that connect to the controller, the AP sends keepalive to both active and standby controllers every 10 minutes. Upon receiving the keepalive, the controller updates the AP's last activity timestamp. If the last activity timestamp is more than 15 minutes, the controller ages out the AP.

High Availability for bridge mode is supported on Campus APs. In this mode, the controller sends ACL Names to the APs instead of the ACL IDs. These APs generate and maintain the mapping between the ACL Name and ACL Id. In the event of a failover the ACL name is sent to the AP from the stand-by controller. Since AP maintains the mapping, the ACL IDs remain intact during a failover.

Redundancy and High Availability Requirements and Limitations

A backup controller can use the High Availability feature. However, a backup controller can only accept standby connections from APs, and will not serve active APs as long as its master redundancy role is **backup**.

This type of High Availability deployment has the following requirements and limitations:

- A backup-master controller can only form an active-standby pair with the master controller.
- The backup master cannot terminate active APs.
- Both the backup-master and master controllers must be configured with the **dual** controller role.

- The **controller <ip>** defined in the high availability group profile must be the IP address of the controller.
- If MultiZone is enabled, High Availability cannot be configured.
- **The inter-controller heartbeat feature is not recommended for backup-master and mastercontroller pairs using the High Availability feature.** If the inter-controller heartbeat feature is enabled in a high availability group profile for redundant masters, the inter-controller failover time must be greater than the VRRP failover time. That is, the (heartbeat interval * heartbeat threshold) value should be greater than the (advertisement time * 3 + preemption delay + skew time [which is based on priority]).

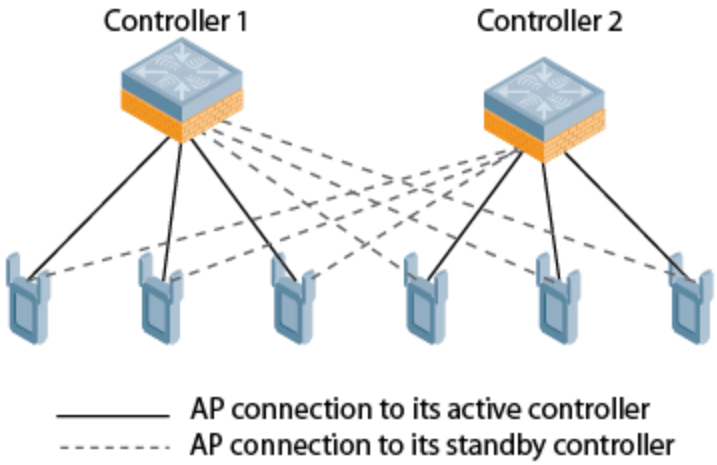
High Availability Deployment Models

High availability supports the following deployment models.

- Active/Active Deployment Model on page 796
- 1:1 Active/Standby Deployment Model on page 796
- N:1 Active/Standby Deployment Model

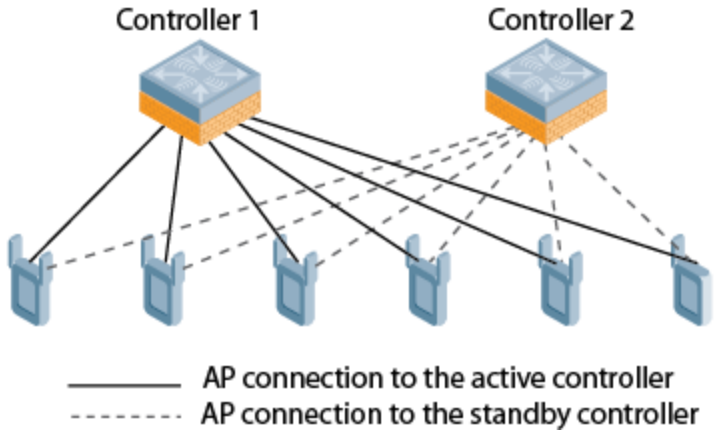
In this model, two controllers are deployed in dual mode. Controller one acts as a standby for the APs served by controller two, and vice-versa. To ensure that each AP gets a standby, Aruba recommends not to have AP count more than 50% of the platform limit; if one controller fails, all the APs served by that controller would failover to the other controller, providing high availability redundancy to all APs in the cluster.

Figure 94 *Active-Active HA Deployment*



In this model, the active controller supports up to 100% of its rated capacity of APs, while the other controller is idle in standby mode. If the active controller fails, all APs served by the active controller will failover to the standby controller.

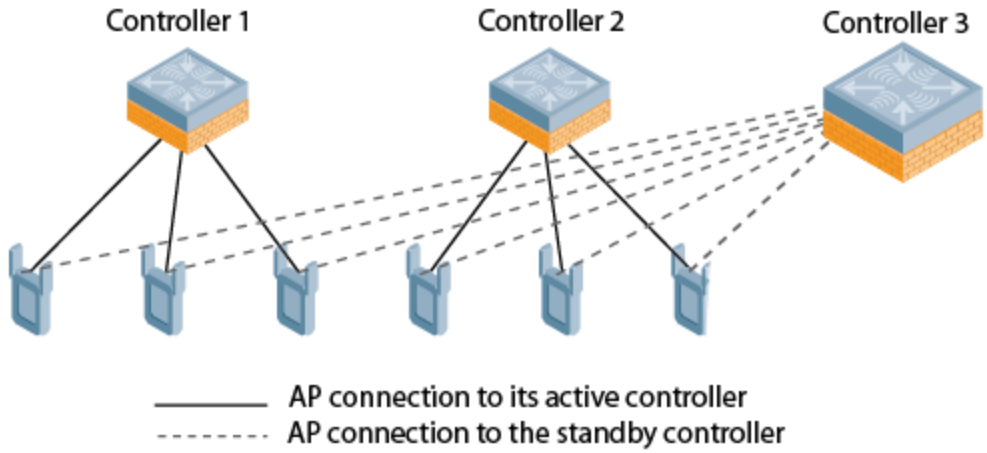
Figure 95 *1:1 Active/Standby Deployment*



In this model, the active controller supports up to 100% of its rated AP capacity, while the other controller is idle in standby mode. If an active controller fails, all APs served by the active controller will failover to the standby controller. This model requires that the AP capacity of the standby controller is able to support the total number of APs distributed across all active controllers in the cluster.

In the cluster shown in the example below, the standby controller has enough AP capacity to support the total number of APs terminating at the active controllers (Controller 1 and Controller 2).

Figure 96 *1:1 Active/Standby Deployment*



High Availability with Extended Capacity

The standby controller over-subscription feature allows a standby controller to support connections to standby APs beyond the controller's original rated AP capacity. The following section of this document gives and lists requirements and capacity limitations for this feature. For more details on enabling the extended standby controller capacity, see Configuring High Availability on page 799.

A controller that acts as a standby controller can oversubscribe to standby APs by up to four times that controller's rated AP capacity, as long as the tunnels consumed by the standby APs do not exceed the maximum tunnel capacity for that standby controller.

Feature Requirements

This feature can be enabled on managed devices where centralized licensing is enabled on the active and standby Mobility Master, or on stand-alone controllers that are not using VRRP-based redundancy.

If centralized licensing is disabled, the standby AP over-subscription feature is also disabled. Standby controller over-subscription and the high availability state synchronization features are mutually exclusive and cannot be enabled simultaneously. If your deployment uses the state synchronization feature, you must disable it before you enable standby controller over-subscription.

Standby Controller Capacity

The following table describes the AP over-subscription capacity maximum supported tunnels and the controllers that support this feature. This feature is not applicable for 70xx controllers.

Table 146: *Controller Support for Standby Oversubscription*

Controller Model	Standby AP Capacity	Maximum Tunnels Supported
7205	4x rated AP capacity	8192 tunnels
7210	4x rated AP capacity	16384 tunnels
7220	4x rated AP capacity	32768 tunnels
7240	4x rated AP capacity	65536 tunnels

To determine the number of standby tunnels consumed by APs on each active controller, multiply the number of APs on the active controllers by the number of BSSIDs per AP. For example, consider a deployment with four active 7210 controllers that each have 512 APs with 8 BSSIDs. The APs on each active controller consume (512 * 8) tunnels, for a combined total of 16,384 tunnels. A single 7210 controller using the standby controller over-subscription feature can act as the standby controller for all four active controllers in this example because this topology is within the 4x rated AP capacity limit and maximum tunnel limit for the 7210 controller model.

If the network administrator later changed all the APs in this deployment to support 10 BSSIDs, each active controller would use (512 * 10) tunnels, for a combined total of 20,480 tunnels on the four active controllers. The tunnels required by the APs on the active controllers would then exceed the maximum tunnel limit for the standby controller, so the standby controller can no longer support all APs on the active controllers. Dynamic changes to configuration (such as the addition of BSSIDs to any AP group) causes all the standby APs to disconnect and reconnect back to the standby controller defined by their updated configuration

To view information about the numbers of currently associated APs and supported BSS tunnels, and the remaining capacity for additional APs and BSS tunnels, issue the **show ha oversubscription statistics** command. For more information refer to the *ArubaOS CLI Reference Guide*.

AP Failover

If a standby controller reaches its AP over-subscription capacity or exceeds its maximum BSSID limit, the standby controller drops any subsequent standby AP connections. A dropped AP attempts to reconnect to the standby controller, but after it exceeds the maximum number of request retries, the AP informs the active controller that it is unable to connect to the standby controller. The active controller then prompts the AP to create a standby tunnel to another standby controller, if one is configured.

If an active controller fails, the APs on the active controller failover to the standby controller. Once the standby controller has reached its capacity for active APs, it terminates tunnels to any standby APs that the controller can no longer serve. When these APs detect that there is no longer a heartbeat between the AP and the standby controller, they notify their active controller that they can no longer connect to the standby. The active controller then prompts the APs to establish standby tunnels to another standby controller, if one is configured.

Client State Synchronization

Client state synchronization allows faster client reauthentication in the event of a controller failure by synchronizing PMK and Key cache entries between active and standby controllers. When you enable this feature, clients only need to perform a four-way key exchange to reconnect to the network (instead of performing a full authentication to the RADIUS server), dramatically shortening the time required for the client to reconnect.



The following section of this document describes topologies, guidelines, and limitations for this feature. To view the procedure for enabling the client state synchronization feature, see Configuring High Availability on page 799.

Feature Guidelines and Limitations

Note the following guidelines and limitations before enabling this feature in your high availability deployment:

- Only APs that support 802.11n and 802.11ac support client state synchronization.
- The client state synchronization and standby controller over-subscription features are mutually exclusive and cannot be enabled simultaneously. If your deployment uses the standby controller over-subscription feature, the feature must be disabled before enabling state synchronization.

High Availability Inter-Controller Heartbeats

The high availability inter-controller heartbeat feature allows for faster AP failover from an active controller to a standby controller, especially in situations where the active controller reboots or loses connectivity to the network.

The inter-controller heartbeat feature works independently from the AP mechanism that sends heartbeats from the AP to the controller. If enabled, the inter-controller heartbeat feature supersedes the AP's heartbeat to its controller. As a result, if a standby controller detects missed inter-controller heartbeats from the active controller, it triggers its standby APs to failover to the standby controller, *even if those APs have not detected any missed heartbeats between the APs and their active controller*. Use this feature with caution in deployments where the active and standby controllers are separated over high-latency WAN links.

When this feature is enabled, the standby controller starts sending regular heartbeats to an AP's active controller as soon as the AP has an UP status on the standby controller. The standby controller initially flags the active controller as **unreachable**, but changes its status to **reachable** as soon as the active controller sends a heartbeat response. If the active controller later becomes unreachable for the number of heartbeats defined by the heartbeat threshold (default of 5 missed heartbeats), the standby controller immediately detects this error and informs the APs using the standby controller to failover from the active controller to the standby controller. If, however, the standby controller never receives an initial heartbeat response from the active controller, and therefore never marks the active controller as initially reachable, the standby controller will not initiate a failover.

This feature is disabled by default. It can be used in conjunction with the high availability state synchronization feature only in topologies that use a single active and standby controller, or a pair of dual-mode active controllers that act as standby controllers for each other. High availability inter-controller heartbeats can be enabled and configured in the high-availability group profile using the WebUI or Command-Line interface.

For more details on how to enable and configure inter-controller heartbeats, see Configuring High Availability on page 799.

Configuring High Availability

The high availability feature supports redundancy models with an active controller pair, or an active/standby deployment model with one backup controller supporting one or more active controllers. Each of these clusters of active and backup controllers comprises a high-availability group. All active and standby controllers within a single high-availability group must be deployed in independent masters topology. An independent masters topology requires all independent master controllers to have the same WLAN configuration.



The High Availability feature is not supported on Standalone Master-Standby deployments.

The following procedure configures High Availability using the WebUI:

1. For a Stand alone controller, under **Mobility Controller**, select your device and then navigate to the **Configuration > Redundancy> High availability** tab.
2. For a managed device, navigate to the **Configuration > Redundancy> High availability** tab.
3. Expand the **HA Groups** accordion and click **+**. A pop-up window appears.
4. In the **Name** field, enter a name for the new HA group.
5. Configure an IPv4 or IPv6 address for the controller.
   a. Click **+** in the **HA Controller IPv4** or **HA Controller IPv6** fields. The **Add HA Controller IP** window opens. Enter an IP address.



IPv4 and IPv6 controllers can be part of the same HA group profile. However, HA works only between controllers of same IP format.

   b. Click the **Role** drop-down list to assign a role to the controller. The IP address of each controller must be reachable by APs and must be the IP address that appears in the **Configuration > Controller > System settings** tab of the controller WebUI, or in the output of the **show controller-ip** CLI command. The role can be one of the following options:

- **active**: Controller is active and serving APs.
- **dual**: Controller serves some APs and acts as a standby controller for other APs.
- **standby**: Controller does not serve APs and only acts as a standby in case of failover.

   c. Click **OK** to add the controller to the group.

6. (Optional) Select the **Pre-emption** check box to enable the failed over APs to attempt to connect back to its original active controller once the controller is reachable again. When you enable this setting, the AP waits for the time specified by the **lms-hold-down-period** parameter defined in the **ap system** profile before the AP attempts to switch back from the standby controller to the original controller.
7. (Optional) The standby controller over-subscription feature allows a standby controller to support connections to standby APs beyond the controller's original rated AP capacity. To enable this feature, select the **Over subscription** check box.
8. (Optional) Select the **State synchronization** check box to enable the feature. State synchronization improves failover performance by synchronizing client authentication state information from the active controller to the standby controller. (For more information about state synchronization, see Client State Synchronization on page 799).



State synchronization is not applicable for IPv6 controllers.

9. (Optional) Select the **Heartbeat** check box to enable the high availability inter-controller heartbeat features, which enable faster AP failover from an active controller to a standby controller, especially in situations where the active controller reboots or loses connectivity to the network. To edit the default heartbeat threshold and interval values:
   - Enter a heartbeat threshold in the **Heartbeat threshold** field to define the number of heartbeats that must be missed before the APs are forced to failover to the standby controller. This value must be between 3 and 10, inclusive.
   - Enter a heartbeat interval in the **Heartbeat interval (ms)** field to define how often inter-controller heartbeats are sent. This value must be between 100 and 1000 ms, inclusive.
10. (Optional) If you enabled the state synchronization feature in Step , enter a pre-shared key into the **Pre-shared key** and **Retype pre-shared key** fields.
11. Click **Submit**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following procedure associates a controller with a High Availability group:

1. For a standalone controller, under **Mobility Controller**, select your device and then navigate to the **Configuration > Redundancy > High availability** tab.
2. For a managed device, navigate to **Configuration > Redundancy > High availability** tab.
3. Expand the **HA Member** accordion.
4. Select a HA group from **Member of HA group** drop-down list.
5. Click **Submit.**
6. Click **Pending Changes.**
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

To configure a High Availability group using the command-line interface, access the CLI in config mode and issue the following commands. The high availability group profile should be configured with a pair of IPv4 controller addresses and pair of IPv6 controller addresses to allow an IPv4 or IPv6 access point to establish a connection to a standby controller.

```
(host) [mynode] (config) #ha group-profile <profile>
(host) [mynode] (HA group information "default") #controller <ip> role
[active/dual/standby]
```

A controller using the high availability features must be defined as a member of a high availability group. To add a controller to the new high availability group, issue following CLI command:

```
(host)[mynode] (config)#ha group-membership <ha-group>
```

Migrating from VRRP or Backup-LMS Redundancy

High Availability: Fast Failover provides redundancy for APs, but not for controllers. Deployments that require master controller redundancy should continue to use an existing VRRP redundancy solution. If your deployment currently uses a backup-LMS or VRRP redundancy solution, use the following procedures to migrate to a High-Availability-based solution. For more information on this topology, see High Availability Deployment Models on page 796.

Migrating from VRRP Redundancy

The following procedure migrates from VRRP to High-Availability redundancy:

1. Remove the VRRP IP address as the LMS IP address of the AP.

```
(host)[mynode](AP system profile) #no lms-ip
```

2. Configure the AP to use the active controller's IP address (not the VRRP IP address) as the LMS-IP for the AP.

```
(host)[mynode](AP system profile) #lms-ip <ipaddress>
```

3. Configure the AP to use the standby controller IP address (not the VRRP IP address) as the backup LMS-IP for the AP.

```
(host)[mynode](AP system profile) #bkup-lms-ip <ipaddress>
```

4. Configure the master controller with a dual role in the high-availability group profile.

```
(host)[mynode](config) #ha group-profile grp1
(host)[mynode](HA group information "grp1"): controller <ipaddress> role dual
```

5. Configure the standby controller with a dual role in the high-availability group profile.

```
(host)[mynode](HA group information "grp1"): controller <ipaddress> role dual
```

Migrating from Backup-LMS Redundancy

The following steps migrates from Backup-LMS to High-Availability redundancy and maintains the existing configuration as defined by the **lms-ip** and **bkup-lms-ip** parameters in the AP system profile.

1. Configure the controller serving the AP with a dual role in the high-availability group profile.

```
(host)[mynode](config) #ha group-profile grp1
(host) (HA group information "grp1"): controller <ipaddress> role dual
```

2. Configure the AP's standby controller with a dual role in the high-availability group profile.

```
(host)[mynode](HA group information "grp1"): controller <ipaddress> role dual
```

VRRP Redundancy



The term controller in this section refers to a stand-alone controller or a managed device running an ArubaOS version 8.x.x.x.

The Virtual Router Redundancy Protocol (VRRP) is used to create various redundancy solutions, such as pairs of controllers acting in active-backup mode or in master-standby mode by using a virtual IP address. 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.

VRRP eliminates a single point of failure by providing a mechanism to elect a VRRP master controller. If VRRP preemption is disabled (the default setting) and all controllers share the same priority, the first controller that comes up becomes the VRRP master. However, if VRRP preemption is enabled and all controllers share the same priority, the controller with the highest IP address becomes the VRRP master.

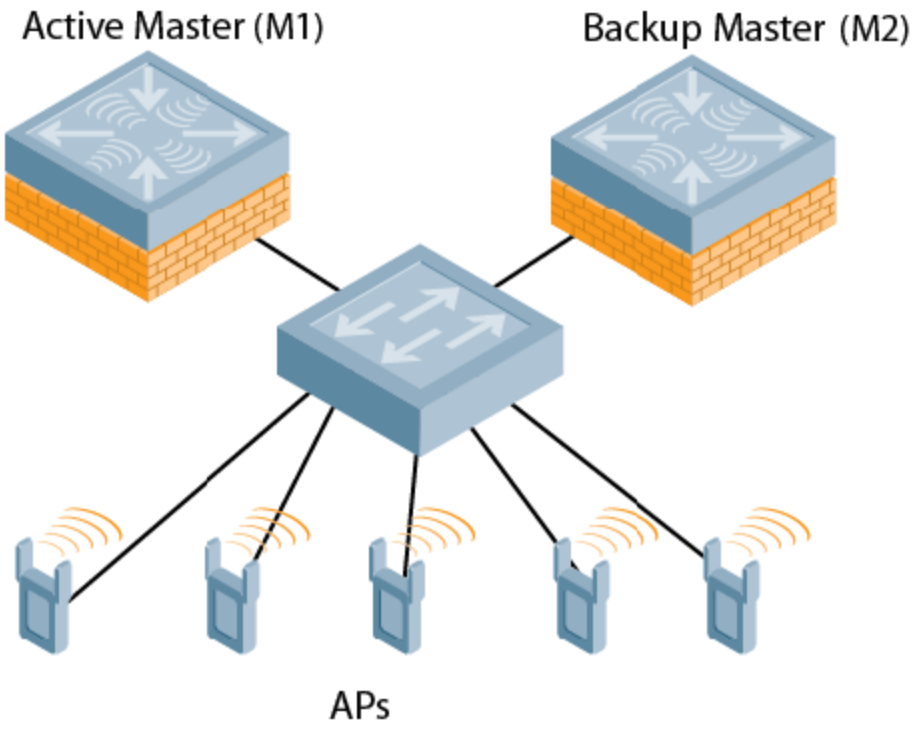
When you need to change the master-ip on the managed device from interface IP of the active Mobility Master to the vrrp-ip of the Mobility Master, the changes need to be performed in the managed device. The reason this needs to be changed in the managed device is because when the master IP is changed on the managed device, the managed device is not aware if the new master-ip is vrrp-ip of the existing Mobility Master or IP of a new Mobility Master.

Therefore, when you change the master-ip, the setup dialog should be executed on the managed device with write erase followed by deleting the device entry on the Mobility Master for this managed device to start setup-dialogue. This is to avoid issues in the network that will be caused by old setup dialogue which will be maintained in the managed device, if we do not do write erase and brought up managed device cleanly.

ArubaOS supports VRRP-based LMS redundancy in a deployment with active-backup redundancy. In the topology illustrated in Figure 97, when an AP connects to the master controller (M1), the AP receives a standby IP. The standby IP is used by the AP to establish a standby connection to the backup master (M2). If the active master becomes unreachable or reboots, the backup master changes its VRRP role to master and accepts active AP connections.

When M1 comes back up, it initially acts as a backup master, and APs associated to M2 establish a standby connection to M1. When the controllers change roles and M1 becomes the active master once again, M2 forces the APs to use M1 as their active master. If an AP has not established a connection to M1 before it disassociates from M2, the AP rebootstraps before it reconnects back to M1.

Figure 97 *Redundancy with an Active-Backup Master Controller Pair*



When a VRRP instance is configured on the controller VLAN, there would be no change in the VRRP state if the failover scenario was tested by shutting down the port or bringing down the VLAN. The controller remains in the Master state and sends VRRP advertisements, which do not reach the peer controller. When the port is down, the peer controller becomes the Master. However, when the port on the previous master is enabled, it takes over the Master state. The peer controller moves out of the master state when the original master sends a higher priority advertisement, even when preemption is not enabled. The peer controller will not be preempted if the master controller crashes or reboots.

Before you Begin

Before you begin configuring VRRP redundancy, obtain the following network information:

- VLAN ID for the two controllers on the same Layer-2 network.
- Virtual IP address to be used for the VRRP instance.

Configuring a Primary and Backup Master for Failover Redundancy

The following procedure configures VRRP on the primary and backup master controllers:

1. For a stand-alone controller, under **Mobility Controller**, select your device and then navigate to the **Configuration > Redundancy > L2 redundancy** tab.
2. For a managed device, navigate to **Configuration > Redundancy > L2 redundancy** in the **Managed Network** node hierarchy.
3. Expand the **Virtual Router Table** accordion.
4. Click **+** to add a new virtual router. The **New Virtual Router** fields appear.
5. Select the IP version from the **IP Version** drop-down list.
6. Select the VLAN on which you want to configure VRRP from the **VLAN** drop-down list.
7. Set **Admin State** to **UP**.
8. Specify the priority value in the **Priority** field. For a backup controller, use the default priority value of 100. For the primary controller, use a priority value higher than the default, such as 110.
9. Configure other VRRP parameters as described in Table 147 .
10. Click **Submit**.
11. Click **Pending Changes**.
12. In the **Pending Changes** window, select the check box and click **Deploy changes**
13. Repeat steps 1-11 to configure VRRP on the other controller in the primary and backup redundant pair.

Ensure to reload the device whenever you modify the Master VRRP configuration under Master Redundancy to avoid any configuration errors.

Table 147: *VRRP Configuration Parameters*

Parameter	Description
ID	The ID uniquely identifies this VRRP instance. For ease in administration, you should configure this with the same value as the VLAN ID.
Description	This is an optional text description to describe the VRRP instance.

Table 147: *VRRP Configuration Parameters*

Parameter	Description
IP version	Select IPv4 \ IPv6 from the drop-down list box.
Authentication Password	This is an optional password of up to eight characters that can authenticate VRRP peers in their advertisements. If this is not configured, there is no authentication password.
Retype authentication password	Reconfirm the password, if configured.
IP address	Based on the selection made in the IP version field, either IP Address \ IPv6 Address is displayed. This is the virtual IP address that will be owned by the elected VRRP master. Ensure that the same IP address and VRRP ID is used on each member of the redundant pair. **NOTE:** The IP address must be unique and cannot be the loopback address of the device. A maximum of only two virtual IPv6 addresses can be configured on each VRRP instance. Only IPv6 address format is supported for the v6 instance.
IPv6 address	Configure the virtual IPv6 address that will be owned by the elected VRRP master. Use the same IPv6 address on each member of the redundant pair. This IPv6 address will be redundant - it will be active on the VRRP master, and will become active on the VRRP backup when the VRRP master fails. **NOTE:** The IPv6 address must be unique and cannot be the loopback address of the device. A maximum of only two virtual IPv6 addresses can be configured on each VRRP instance. Starting from ArubaOS 8.2.1.0, you can configure a unique local address as the VRRP IPv6 address on the Mobility Master and the managed devices.
Priority	Priority level of the VRRP instance for the device. This value is used in the election mechanism for the master. When configuring VRRP on a standby device, use the default priority value of 100. For a master device, use a higher priority value, such as 110.
Advertisement interval (secs)	This is the interval, in seconds, between successive VRRP advertisements sent by the current master. The default interval time is recommended. Default: 1 second
Enable router Pre-emption	Selecting this option means that a device can take over the role of master if it detects a lower priority device currently acting as master.
Pre-emption delay (secs)	Specifying a value enables the delay timer. The timer is triggered when the VRRP state moves out of backup or init state to become a master. This is applicable only if you enable router pre-emption. When the timer is triggered, it forces VRRP to wait for a specified period of time, so that all the applications are ready before coming up. This prevents the APs from connecting to the managed device or the stand-alone controller before it can receive them. In the meantime, if there is an advertisement from another VRRP, the VRRP stops the timer and does not transition to master.
Hold Time	Hold time is the number of seconds until which a managed device waits, before starting VRRP to account for System or Network convergence delays. Default: 45 seconds **NOTE:** Configuring the hold time will not take affect if preemption is enabled.

Table 147: *VRRP Configuration Parameters*

Parameter	Description
Admin state	Administrative state of the VRRP instance. To start the VRRP instance, change the admin state to **UP** in the WebUI.
VLAN	VLAN on which the VRRP protocol runs.
Tracking master up-time	(Optional) Perform VRRP priority tracking based on how long the device has been the master. This feature is designed to ensure that a master will only be allowed to take and maintain control of the VRRP if it has been up for a certain amount of minutes (0-1440). This prevents an issue where a device that is periodically going up and down assumes the role of primary master.
Tracking master up-time priority	(Optional) The additional priority given to the master once it has been up for the time interval defined by the **Tracking Master Up-time** parameter.
Tracking VRRP master state ID	(Optional) Perform tracking based on the UP or DOWN state of another VRRP master by specifying the VRRP ID of the master to be tracked.
Tracking VRRP master state priority	(Optional) The priority taken away from a VRRP master if it is in a DOWN state. The priority levels are returned to their previous state when the VRRP master comes back up.
Tracking VLAN	(Optional) Perform VRRP priority tracking based on the UP or DOWN state of a VLAN. Click **+** below the **Tracking VLAN** table and specify the following values: ■ VLAN Id: ID of the VLAN to be tracked. ■ Subtract: Priority level to be subtracted from the device's VRRP priority if the tracked VLAN goes down.
Tracking interface	(Optional) Perform VRRP priority tracking based on the UP or DOWN state of a specific interface. Click **+** below the **Tracking Interface** table and specify the following values: ■ Interface: Interface Port to be tracked. ■ Subtract: Priority level to be subtracted from the device's VRRP priority if the tracked interface goes down.

Execute the following commands to configure a new virtual router on a stand-alone controller:

```
(host) [mynode] (config) #vrrp <id>
(host) [mynode] (config-submode)#ip address <ip-address>
(host) [mynode] (config-submode)#vlan <vlanID>
(host) [mynode] (config-submode)#priority <0-255>
```

Execute the following commands to configure a new virtual router on a managed device:

```
(host) [md] (config) #vrrp <id>
(host) [md] (config-submode)#ip address <ip-address>
(host) [md] (config-submode)#vlan <vlanID>
(host) [md] (config-submode)#priority <0-255>
```

Configuring APs to use the VRRP IP

Configure the APs associated with the master controller to terminate their tunnels on the VRRP 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.



This configuration must be executed on the master controller; the APs obtain their configuration from the master controller.

The following procedure configures VRRP on an AP system profile:

1. For a stand-alone controller, under **Mobility Controller**, select your device and then navigate to **Configuration > System > Profiles** tab.
2. For a managed device, navigate to **Configuration > System > Profiles** in the **Managed Network** node hierarchy.
3. Under **All Profiles > AP**, expand **AP system**.
4. Select the AP system profile for which you want to configure VRRP.
5. Expand the **LMS Settings** accordion and enter the virtual IP address into the **LMS IP** field.
6. Click **Save**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The follow procedure configures VRRP for an AP group:

1. For standalone controller, under **Mobility Controller**, select your device and then navigate to **Configuration > AP Groups**.
2. Select the **LMS** tab from the selected AP group table.
3. Enter the virtual IP address into the **IP address** field. For IPv6 address, enter the value in the **IPv6 address** field.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure VRRP on an AP system profile and applies it to an AP profile and an AP group:

On a stand-alone controller:

```
(host) [mynode] (config) #ap system-profile <profile-name>
(host) [mynode] (AP system profile "<profile-name>") #lms-ip <ip-address>
(host) [mynode] (AP system profile "<profile-name>") #ap-name <ap-profile-name>
(host) [mynode] (AP name "<ap-profile-name>") #ap-system-profile <profile-name>
(host) [mynode] (AP name "<ap-profile-name>") #exit
(host) [mynode] (config) #ap-group <ap-group-name>
(host) [mynode] (AP group "<ap-group-name>") #ap-system-profile <profile-name>
```

On a managed device:

```
(host) [md] (config) #ap system-profile <profile-name>
(host) [md] (AP system profile "<profile-name>") #lms-ip <ip-address>
(host) [md] (AP system profile "<profile-name>") #ap-name <ap-profile-name>
(host) [md] (AP name "<ap-profile-name>") #ap-system-profile <profile-name>
(host) [md] (AP name "<ap-profile-name>") #exit
(host) [md] (config) #ap-group <ap-group-name>
(host) [md] (AP group "<ap-group-name>") #ap-system-profile <profile-name>
```

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 Master Redundancy and Database Synchronization

In a redundant master controller scenario, you can configure a redundant pair to synchronize their WMS and local user databases. You can either manually or automatically synchronize the databases. When manually synchronizing the database, the active VRRP master synchronizes its database with the

standby. The command takes effect immediately. When configuring automatic synchronization, you set how often the two controllers synchronize their databases. To ensure successful synchronization of database events, you should set periodic synchronization to a minimum period of 20 minutes.



The master-standby configuration and the database synchronization between the redundant controllers is not applicable to controllers configured as managed devices. However, it is applicable to stand-alone controllers running ArubaOS 8.x.x.x.

The following procedure configures master redundancy:

1. For standalone controller, under **Mobility Controller**, select your device and then navigate to the **Configuration > Redundancy > L2 redundancy** tab.
2. In the Mobility master node hierarchy, navigate to the **Configuration > Redundancy > L2 redundancy** tab.
3. Under **Master Redundancy**, do the following:
4. Enter the VRRP ID to be associated to the master-redundancy pair in the **Master VRRP** field.
5. Enter the IP address of the redundancy pair in the **IP address of peer** field.
6. Select the authentication method from the **Authentication** drop-box.
    a. If **IPSec key** is selected as an authentication method, enter the **IPSec key of the peer** and **Re-type the key**
    b. If **Certificate** is selected as an authentication method and **Factory** is selected as the **Certificate type,** enter the**Peer's MAC address**
    c. If **Certificate** is selected as an authentication method and **Custom** is selected as the certificate type, enter the **Peer's MAC address, CA certificate, Server Certificate** and select a **Suite B algorithm** from the drop-down list.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The follow CLI commands configure a master redundancy pair. This configuration is applicable on both active and the standby controllers:

```
(host) [mynode] (config) #master-redundancy
(host) [mynode] (config-submode)#master-vrrp <vrrp-id>
(host) [mynode] (config-submode)#peer-ip-address <ip-address>
(host) [mynode] (config-submode)#write memory
```

The follow CLI commands configure synchronization:

```
(host) [mynode] (config) #database synchronize period
```

To view the database synchronization settings on the controller, use the following command:

```
(host) [mynode] #show database synchronize
```

Chapter 30
IP Mobility

A mobility domain is a group of Aruba managed devices among which wireless users can roam without losing their IP address. Mobility domains are not tied with the Mobility Master; thus, it is possible for a user to roam between managed devices as long as all the managed devices belong to the same Mobility Master.

You enable and configure mobility domains only on Aruba managed devices. No additional software or configuration is required on wireless clients to allow roaming within the domain.

Topics in this chapter include:

- Understanding Aruba Mobility Architecture on page 809
- Configuring Mobility Domains on page 810
- Tracking Mobile Users on page 812
- Configuring Advanced Mobility Functions on page 814
- Understanding Bridge Mode Mobility Deployments on page 822
- Monitoring Network Traffic Using IP Flow Information Export on page 823
- Enabling Mobility Multicast on page 826

Understanding Aruba Mobility Architecture

Aruba'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 Aruba mobility solution does not require that you install mobility software or perform additional configuration on wireless clients. The Aruba managed devices 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 Aruba managed device in the foreign network with which the mobile client is associated.

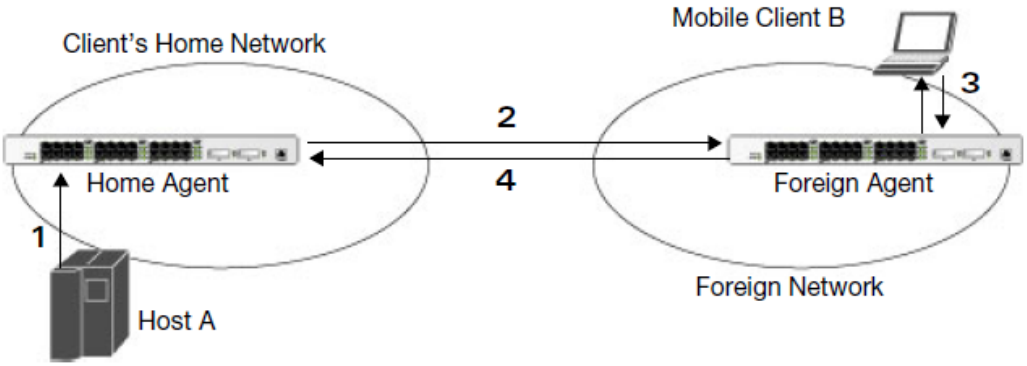
The home agent for the client is a managed device at which the client appears for the first time upon joining the mobility domain. The home agent is the single point of contact for the client when the client roams. The foreign agent for the client is the managed device which handles all Mobile IP communication 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 98 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 Aruba managed device in the foreign network.

The numbers in the Figure 98 correspond to the following descriptions:

1. Traffic to Mobile Client B arrives at the client's home network over 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 98 *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 Managed devices that support the VLANs into which employee users can be placed should be part of the same mobility domain.

NOTE

Aruba mobility domains are supported only on Arubamanaged devices.

A managed device can be part of multiple mobility domains, although it is recommended that a managed device belong to only one domain. The managed device in a mobility domain do not need to be managed by the same Mobility Master.

You configure a mobility domain on a Mobility Master; the mobility domain information is pushed to all managed device that are managed by the Mobility Master. On each managed device, you must specify the active domain (the domain to which the managed device belongs). If you do not specify the active domain, the managed device will be assigned to a predefined default domain.

The basic tasks you need to perform to configure a mobility domain are listed below. The sections following describe each task in further detail.

Table 148: *Tasks to Configure a Mobility Domain*

On a Mobility Master:	On all managed devices in the mobility domain:
Configure the mobility domain, including the entries in the home agent table	Enable mobility (disabled by default) Join a specified mobility domain (not required for default mobility domain)

You can enable or disable IP mobility in a virtual AP profile (IP mobility is enabled by default). When you enable IP mobility in a virtual AP profile, the ESSID that is configured for the virtual AP supports layer-3 mobility. If you disable IP mobility for a virtual AP, any clients that associate to the virtual AP will not have mobility service.

Listed below are the steps to configure and join a mobility domain:

Configuring a Mobility Domain

You configure mobility domains on Mobility Master. All managed devices managed by the Mobility Master share the list of mobility domains configured on the Mobility Master. Mobility is disabled by default and must be explicitly enabled on all managed devices that will support client mobility. Disabling mobility does not delete any mobility related configuration.

The home agent table maps a user VLAN IP subnet to potential home agent addresses. When you enable mobility the managed device to which the client connects for the first time becomes its home agent. The mobility feature uses the home agent 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 managed device with direct access to the user VLAN IP subnet. (Some network topologies may require multiple home agents.)

It is recommended that you configure the switch IP address to match the AP's managed device, or define the VRRP IP address to match the VRRP IP used for managed device redundancy. Do not configure both switch IP address and VRRP IP address as the home agent address, or multiple home agent discoveries may be sent to the managed device.



All user VLANs that are part of a mobility domain must have an IP address that can correctly forward layer-3 broadcast multicast traffic to clients when they are away from the home network.

The mobility domain named *default* is the default active domain for all managed devices. 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 managed device to a user-defined domain, it automatically leaves the default mobility domain. If you want a managed device 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 managed device.

The following procedure describes how to configure a mobility domain:

1. In a **Managed Network** node hierarchy, navigate to **Configuration** > **Services** page and select the **IP Mobility** tab.
2. Expand **Mobility Domain** accordion.
3. To configure the default mobility domain, select the default domain in the **IP Mobility Configuration** table.
4. To create a new mobility domain, click **+** in the **IP Mobility Configuration** table. Enter the value of the **Name** and **Description** fields in the **Create IP Mobility** table.
5. Click **Submit**.
6. Select the newly-created domain name and click +in the **IP Mobility Configuration** table. The **Home Agent** table is displayed.
7. Click **+** in the **Home Agent** table. A **Create Home Agent** table is displayed.
8. Enter the value of the **IP** and **Description** fields in the **Create Home Agent** table.
9. Click **Submit**.

10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure a mobility domain:

```
(host) [md] (config)#router mobile
(host) [md] (config)#ip mobile domain <name>
(host) [md] (config-submode)#hat <home-agent> description <dscr>
```

To view currently-configured mobility domains in the CLI, use the `show ip mobile domain` command.

Ensure 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 managed device to a specific mobility domain is the key to defining the roaming area for mobile clients. You should take extra care in planning your mobility domains and survey the user VLANs and managed device to which clients can roam, to ensure that there are no roaming holes.

All managed device are initially part of the default mobility domain. If you use the default mobility domain, you do not need to specify this domain as the active domain on a managed device. However, once you assign a managed device to a user-defined domain, the default mobility domain is no longer an active domain on the managed device.

The following procedure describes how to activate a mobility domain:

1. In a **Managed Network** node hierarchy, select a device navigate to **Configuration** > **Services** page and select the **IP Mobility** tab.
2. Expand **Mobility Domain** accordion. Click **Enable IP mobility** check box.
3. Select a Domain Name from the **IP Mobility Configuration** table.
4. Select the **Active** check box.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command activates a mobility domain:

```
(host) [md] (config)#ip mobile active-domain <name>
```

To view the active domains in the CLI, use the **show ip mobile active-domains** command on the managed device.

Tracking Mobile Users

This section describes how you can view information about the status of mobile clients in the mobility domain.

Location-related information for users, such as roaming status, AP name, ESSID, BSSID, and physical type are consistent in both the home agent and foreign agent. The username, role, and authentication can be different on the home agent and foreign agent, as explained by the following:

L2 GRE tunnels are automatically established between managed devices in mobility domain at the time of boot up. Whenever a client connects to a managed device in a mobility domain, layer-2 authentication is performed and the station obtains the layer-2 (logon) role. When the client roams to other networks, layer-2 authentication is performed and the client obtains the layer-2 role. If layer-3

authentication is required, this authentication is performed on the client's home agent only. The home agent obtains a new role for the client after layer-3 authentication; this new role appears in the user status on the home agent only. Even if reauthentication occurs after the station moves to a foreign agent, the display on the foreign agent still shows the layer-2 role for the user.

Listed below are the steps to view mobile client and user roaming status, mobile client roaming locations, and home agent discovery on association:

Mobile Client Roaming Status

You can view the list of mobile clients and their roaming status on any managed device in the mobility domain:

```
(host) [md] #show ip mobile host
```

Roaming status can be one of the following:

Table 149: *Client Roaming Status*

Roaming Status Type	Description
Home Switch/Home VLAN	This managed device 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 managed device is not the home agent for a client.
Mobile IP Binding (away)	This managed device is the home agent for a client that is currently away.
Home Switch/Foreign VLAN	This managed device 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 managed device has received a disassociation message for a client, but has not received an association or registration request for the client from another managed device, or a home agent binding for the station has expired before being refreshed by a foreign agent.
No Mobility Service	The managed device cannot provide mobility service to this client. The mobile client may lose its IP address if it obtains the address over DHCP and has limited connectivity. The mobile client may be using an IP address that cannot be served, or there may be a roaming hole due to improper configuration.

Viewing User Roaming Status

The following CLI command displays the roaming status of users on any managed device in the mobility domain:

```
(host) [md] #show user
```

Roaming status can be one of the following:

Table 150: *User Roaming status*

Status Type	Description
Wireless	This client is on its home agent managed device and the client is on the VLAN on which it has an IP address.
Visitor	This client is visiting this managed device and the managed device is not its home agent.
Away	This client is currently away from its home agent managed device.
Foreign VLAN	This client is on its home agent managed device 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.

The following CLI command traces the local/remote host:

```
#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 using the following command:

```
(host) [md] #show ip mobile trail <ip-address>|<mac-address>
```

Home Agent Discovery on Association

In normal circumstances, a managed device performs a home agent discovery only when it is aware of the client's IP address which it learns through the ARP or any Layer-3 packet from the client. This limitation of learning the client's IP and then performing the home agent 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 hand-held devices, Wi-Fi phones and so on. This delays home agent discovery and eventually results in any loss of downstream traffic that is meant for the mobile client.

When home agent discovery on association is triggered, the foreign agent managed device to which the client is associated, sends a unicast request to all managed device within the mobility domain to find if any one of the managed device has the IP mobility state information of the client.

With home agent discovery on association, a managed device can perform a home agent discovery as soon as the client is associated. This feature is enabled by default. This option will also poll for all potential home agents.

The following CLI command configures the mobility association:

```
(host) [md] (config)#wlan virtual-ap default ha-disc-onassoc
```

Configuring Advanced Mobility Functions

Listed below are the key mobility functions, important points to remember, and sample configurations.

Mobility Functions

The following procedure describes how to configure various parameters that pertain to mobility functions on a managed device in a mobility domain:

1. In a **Managed Network** node hierarchy, navigate to **Configuration** > **Services** and select the **IP Mobility** tab.
2. Expand **Global Parameters** accordion and configure the IP mobility settings.

Table 151: *IP Mobility - Global Parameters*

Parameter	Description
Foreign Agent	
Lifetime	Requested lifetime, in seconds, as per RFC 3344, IP Mobility Support for IPv4. Range: 40-65534 seconds Default: 40 seconds
Max. visitors allowed	Set a maximum allowed number of active visitors. Range: 0-5000 visitors Default: 5000 visitors
Registration requests retransmits	Maximum number of times the foreign agent attempts mobile IP registration message exchanges before giving up. Range: 0-5 attempts Default: 3 attempts
Registration requests interval	Retransmission interval, in milliseconds. Range: 100-10000 milliseconds Default: 1000 milliseconds
Home Agent	
Replay	Time difference, in seconds, for timestamp-based replay protection, as described by RFC 3344, IP Mobility Support for IPv4. 0 disables replay. Range: 0-5000 seconds Default: 5000 seconds.
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 managed device, which will become its home managed device. Range: 0-300 seconds Default: 7 seconds
Proxy Mobile IP	
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 the client roams to a different VLAN or managed device.
Mobility trail logging	Enables logging at the notification level for mobile client moves.
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. Valid range is 30 - 300, and by default the value is set to 180.

Parameter	Description
Max. station mobility events per second	Maximum number of mobility events (events that can trigger mobility) handled per second. Mobility events above this threshold are ignored. This helps to control frequent mobility state changes when the client bounces back and forth on APs before settling down. Range: 1-100 events Default: 10 events
Station trail timeout	Specifies the maximum interval, in seconds, an inactive mobility trail is held. Range: 120-3600 seconds Default: 600 seconds
Station trail max. entries	Specifies the maximum number of entries (client moves) stored in the user mobility trail. Range: 1-30 entries Default: 10 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 managed device. The default is 60 seconds but can be safely increased. In many case a station state is deleted without waiting for the stale timeout; user delete from management, foreign agent to foreign agent handoff, and so on. (This is different from the no-service-timeout; no-service-timeout occurs up front, while the stale-timeout begins when mobility service is provided but the connection is disrupted for some reason). Range: 30-3600 seconds Default: 60 seconds
Revocation	
Retransmits	Maximum number of times the home agent or foreign agent attempts mobile IP registration or revocation message exchanges before giving up. Range: 0-5 retransmissions Default: 3 retransmissions.
Interval	Retransmission interval, in milliseconds. Range: 100-10000 milliseconds Default: 1000 milliseconds

3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures foreign agent functionality:

```
(host) [md] (config)#ip mobile foreign-agent {lifetime <seconds> | max-visitors <number>
| registrations {interval <msecs> | retransmits <number>}}
```

The following CLI command configures configure home agent functionality:

```
(host) [md] (config)#ip mobile home-agent {max-bindings <number>|replay <seconds>}
```

The following CLI command configures proxy mobile IP and DHCP functionality:

```
(host) [md] (config)#ip mobile proxy auth-sta-roam-only | block-dhcp-release | event-
threshold <number> | log-trail | no-service-timeout <seconds> | on-association | refresh-
stale-ip | stale-timeout <seconds> | stand-alone-AP | trail-length <number> |trail-
timeout <seconds >
```

The following CLI command configures revocation functionality:

```
(host) [md] (config)#ip mobile revocation {interval <msec>|retransmits <number>
```

The following CLI command enables packet trace for a given MAC address:

```
(host) [md] (config)#ip mobile packet-trace <host MAC address>
```

Proxy Mobile IP

The proxy mobile IP module in a mobility-enabled managed device 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 home agent table using the mobile client's IP address. If there is more than one possible home agent for a mobile client in the home agent table, 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 managed device, it is recommended that you keep the **on-association** option enabled. This helps trigger mobility as soon as 802.11 association packets are received from the mobile client.

Revocations

A home agent or foreign agent can send a registration revocation message, which revokes registration service for the mobile client. For example, when a mobile client roams from one foreign agent to another, the home agent can send a registration revocation message to the first foreign agent so that the foreign agent can free any resources held for the client.

IPv6 L3 Mobility

ArubaOS supports IPv6 L3 Mobility functionality. The existing L3 mobility solution has been enhanced to support dual stacked (IPv4 and IPv6) and pure IPv6 mobile clients. The IPv6 L3 mobility allows the wireless clients to retain their IPv4 or IPv6 addresses across different VLANs within a managed device and between different managed devices. In the previous release, the Aruba Managed devices supported L3 mobility only for single stacked IPv4 clients.

The following changes in the existing behavior is observed in the Aruba managed device when you enable IPv6 L3 Mobility support :

- The managed device throttles and proxies RAs if the router mobile command is enabled.
  - The following command configures the maximum time allowed between sending unsolicited multicast router advertisements from each interface when RA proxy is enabled:
  - (host) [mynode] (config)# ipv6 proxy-ra interval <180-1800>
  - The default value for `proxy-ra` interval is 600 seconds. If RA is configured on an external router, but not within the managed device, the managed device stores the RA in cache and replays the RA from the external server and replays them every proxy-ra interval. If RA is configured in both an external router and in the managed device, clients serviced by the managed device receive RA only from the managed device and not from the external router.
- Layer-3 mobility support for wired and third-party APs are deprecated.
- The HA discovery on association parameter is turned on by default and is not configurable.



By enabling Layer-3 mobility feature, both the solicited RAs and the unsolicited periodic RAs will be converted to Layer-2 unicast and sent to the WLAN clients.

It is recommended to reboot the managed device when you issue the **no router mobile** command so that mutlicast RAs do not continue to get converted to unicast RAs.

Multicast Mobility

Multicast mobility ensures a client gets an uninterrupted multicast stream while roaming. ArubaOS provides support for a MLD proxy to enable IPv6 multicast mobility. To achieve multicast mobility, the Home Agent and the Foreign Agent must be capable of MLD proxying by exchanging the MLD membership information and process MLD messages. ArubaOSmanaged device supports MLD versions v1 and v2.

Important Points to Remember

- ArubaOS does not support the source-based forwarding functionality of MLDv2.
- The multicast traffic flow stops for few seconds for roaming clients after enabling or disabling the DMO option.

The following CLI commands enable MLD proxy in the VLAN:

```
(host)[md](config)# interface vlan <vlan-id>
(host)[md](config-subif)# ipv6 mld proxy <gigabitethernet/fastethernet>
<slot/module/port>
```

The following CLI command displays the interface-specific MLD proxy group information:

```
(host) [md] #show ipv6 mld proxy-group
```

The following CLI command displays the MLD proxy mobility database group information for tracking:

```
(host) [md] #show ipv6 mld proxy-mobility-group
```

The following CLI command displays the statistics of the MLD proxy:

```
(host) [md] #show ipv6 mld proxy-stats
```

The following CLI command displays the MLD proxy mobility multicast statistics:

```
(host) [md]# show ipv6 mld proxy-mobility-stats
```

The following CLI command displays the discovery count table that keeps track of per client home agent discovery:

```
(host) [md] #show datapath mobility discovery-table
```

The following CLI command displays the datapath HA table information:

```
(host) [md] #show datapath mobility home-agent-table
```

The following CLI command displays the mobility multicast-group table that floods the multicast RA traffic to the roaming clients:

```
(host) [md] #show datapath mobility mcast-table
```

The following CLI command displays the statistics of the datapath mobility:

```
(host) [md] #show datapath mobility stats
```

The following CLI command displays the mobility multicast VLAN table information:

```
(host) [md] #show ip mobile multicast-vlan-table
```

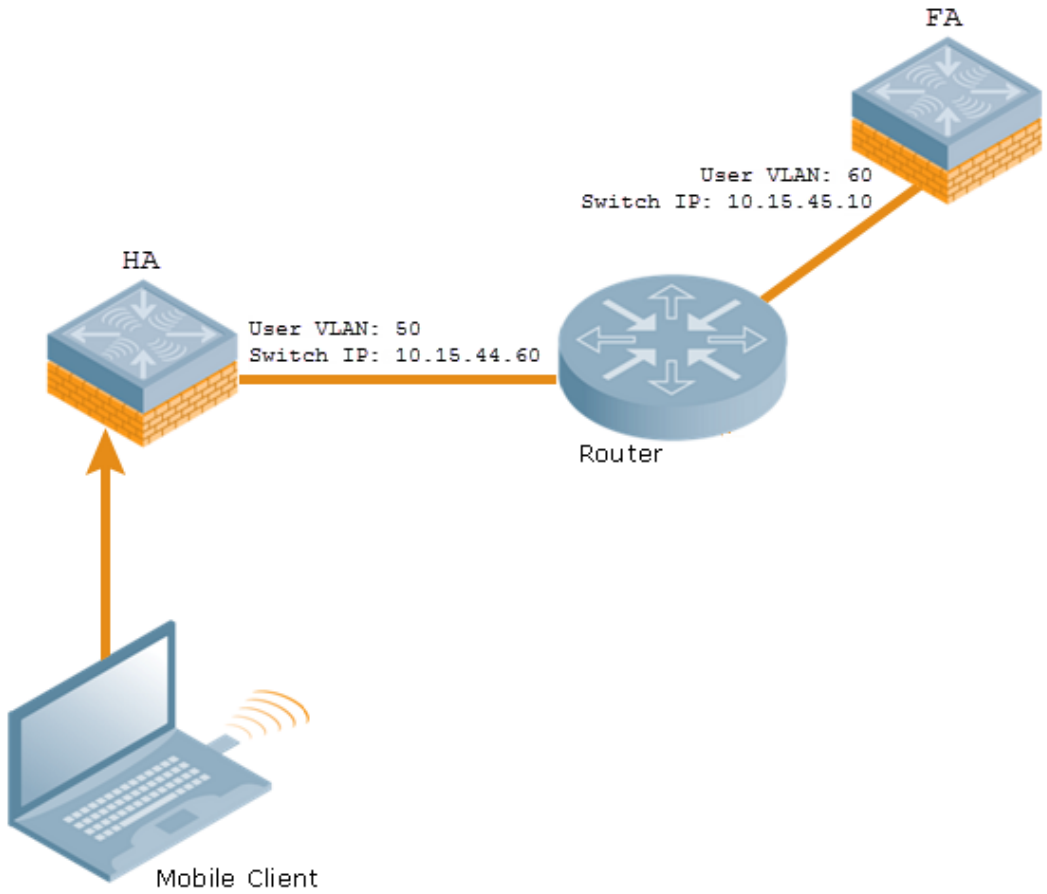
The outputs of the following commands are enhanced to support IPv6 Layer-3 mobility:

- `show ip mobile host`
- `show ip mobile trace`
- `show ip mobile remote`
- `show ip mobile binding`
- `show ip mobile visitor`
- `show ip mobile trail`
- `show ip mobile packet-trace`
- `clear ip mobile trail <IPv6_addr>`
- `show ip mobile traffic`
- `show ip mobile global`
- `show ip mobile hat`
- `show ip mobile domain`
- `ip mobile domain <name> hat <home-agent> description <dscr>`

Sample Configuration

The following figure provides information on how a client moves from one managed device to another, when you enable IPv6 Layer-3 mobility feature:

Figure 99 *Sample IPv6 Layer-3 Mobility Configuration*



The following CLI command displays the initial configuration on the home agent and the foreign agent:

```
(host-HA) #show ip mobile domain
Mobility Domains:, 2 domain(s)
------------------------------
Domain name default
Home Agent Table
Domain name 6.3mobility
Home Agent Table
Home Agent Description
--------------- ----------------
10.15.45.10
10.15.44.60
(host-FA) #show ip mobile domain
Mobility Domains:, 2 domain(s)
------------------------------
Domain name default
Home Agent Table
Domain name 6.3mobility
Home Agent Table
Home Agent Description
--------------- ----------------
```

```
10.15.45.10
10.15.44.60
```

The following CLI command displays information on the client association to a home agent:

```
(host-HA) #show user
Users
-----
IP MAC Name Role Age(d:h:m) Auth VPN l
ink AP name Roaming Essid/Bssid/Phy Profile Forward mode
Type Host Name
---------- ------------ ------ ---- ---------- ---- -
------- ------- ------- --------------- ------- --------
---- ---- -
50.50.50.11 24:77:03:9e:dc:4c authenticated 00:00:00
AP124-B11550-Jibin Wireless mobility-test/00:1a:1e:82:b3:10/a-HT default tunnel
2001:5000::2677:3ff:fe9e:dc4c 24:77:03:9e:dc:4c authenticated 00:00:00
AP124-B11550-Jibin Wireless mobility-test/00:1a:1e:82:b3:10/a-HT default tunnel
fe80::2677:3ff:fe9e:dc4c 24:77:03:9e:dc:4c authenticated 00:00:00
AP124-B11550-Jibin Wireless mobility-test/00:1a:1e:82:b3:10/a-HT default tunnel

(host-HA) #show ip mobile host
Mobile Host List, 1 host(s)
---------------------------
24:77:03:9e:dc:4c
IPv4: 50.50.50.11
IPv6: fe80::2677:3ff:fe9e:dc4c, 2001:5000::2677:3ff:fe9e:dc4c
Roaming Status: Home Switch/Home VLAN, Service time 0 days 00:00:57
Home VLAN 50

(host-HA) #show datapath bridge table 24:77:03:9e:dc:4c
Datapath Bridge Table Entries
-----------------------------
Flags: P - Permanent, D - Deny, R - Roamed Client, M - Mobile, X - Xsec, A - Auth, O -
Outer V
LAN, T - Trusted
MAC VLAN Assigned VLAN QinQ VLAN Destination Flags
----------------- ---- ------------- --------- ----------- -----
24:77:03:9E:DC:4C 50 50 0 tunnel 17 PM

(host-HA) #show datapath station
Datapath Station Table Entries
------------------------------
Flags: W - WEP, T - TKIP, A - AESCCM, M - WMM N - .11n client
S - AMSDU, G - AESGCM, R - DATA READY, I - INACTIVE, r - ROAMED
MAC BSSID VLAN Bad Decrypts Bad Encrypts Cpu Qsz RSN cap Aid HomeVlan
Flags
----------------- ----------------- ---- ------------ ------------ --- -------- ------- -
--- -
------- -----
24:77:03:9E:DC:4C 00:1A:1E:82:B3:10 50 0 0 8 0 0 0 0 0000 0001
50 MN
```

The following CLI command displays the status of a client roaming to a foreign agent:

```
(host-FA) #show ap association
Association Table
-----------------
Name bssid mac auth assoc aid l-int essid vlan-i
d tunnel-id phy assoc. time num assoc Flags Band steer moves (T/S)
---- ----- --- ---- ----- --- ----- ----- ------
- --------- --- ----------- --------- ----- ----------------------
Ap_local 6c:f3:7f:3a:ba:d8 24:77:03:9e:dc:4c y y 1 100 mobility-test 60
0x1000f a-HT-40sgi-2ss 3m:20s 1 WA 0/0
```

```
Num Clients:1

(host-FA) #show user
Users
-----
IP MAC Name Role Age(d:h:m) Aut
h VPN link AP name Roaming Essid/Bssid/Phy Profile Forward mode T
ype Host Name
---------- ------------ ------ ---- ----------
---- -------- ------- ------- --------------- ------- ------------
---- --
50.50.50.11 24:77:03:9e:dc:4c sys_mip_role_649130_9 00:00:03
Ap_local Visitor mobility-test/6c:f3:7f:3a:ba:d8/a-HT default tunnel
Win 7
2001:5000::2677:3ff:fe9e:dc4c 24:77:03:9e:dc:4c sys_mip_role_649130_9 00:00:03
Ap_local Visitor mobility-test/6c:f3:7f:3a:ba:d8/a-HT default tunnel
Win 7
User Entries: 2/2
Curr/Cum Alloc:1/7 Free:1/6 Dyn:2 AllocErr:0 FreeErr:0
(host-FA) #show ip mobile host
Mobile Host List, 1 host(s)
---------------------------
24:77:03:9e:dc:4c
IPv4: 50.50.50.11
IPv6: 2001:5000::2677:3ff:fe9e:dc4c
Roaming Status: Mobile IP Visitor, Service time 0 days 00:03:33
Home VLAN 50, visiting local VLAN 60

(host-FA) #show datapath bridge table 24:77:03:9e:dc:4c
Datapath Bridge Table Entries
-----------------------------
Flags: P - Permanent, D - Deny, R - Roamed Client, M - Mobile, X - Xsec, A - Auth, O -
Outer V
LAN, T - Trusted
MAC VLAN Assigned VLAN QinQ VLAN Destination Flags
----------------- ---- ------------ --------- ----------- -----
24:77:03:9E:DC:4C 4095 60 0 tunnel 15 PMR
24:77:03:9E:DC:4C 60 60 0 tunnel 15 PM

(host-FA) #show datapath station
Datapath Station Table Entries
------------------------------
Flags: W - WEP, T - TKIP, A - AESCCM, M - WMM N - .11n client
S - AMSDU, G - AESGCM, R - DATA READY, I - INACTIVE, r - ROAMED
MAC BSSID VLAN Bad Decrypts Bad Encrypts Cpu Qsz RSN cap Aid HomeVlan
Flags
----------------- ----------------- ---- ------------ ------------ --- -------- ------- -
--- -
------- -----
24:77:03:9E:DC:4C 6C:F3:7F:3A:BA:D8 60 0 0 7 0 0 0 0 0000 0001
50 MNr

(host-FA) #show ip mobile visitor
Foreign Agent Visitor list, 1 host(s)
-------------------------------------
24:77:03:9e:dc:4c
IPv4: 50.50.50.11
IPv6: 2001:5000::2677:3ff:fe9e:dc4c
HA Addr 10.15.44.60, Registration id D51BA8BC:856865FC
Lifetime granted 00:00:40 (40), remaining 00:00:36
Tunnel id 9, src 10.15.44.10 dest 10.15.44.60, reverse-allowed
```

The following CLI command displays the status of the client on a home agent after roaming:

```
(host-HA) #show user
Users
-----
IP MAC Name Role Age(d:h:m) Auth VPN l
ink AP name Roaming Essid/Bssid/Phy Profile Forward mode Type Hos
t Name
---------- ------------ ------ ---- ---------- ---- -
------- ------- ------- --------------- ------- ------------ ----
---------
50.50.50.11 24:77:03:9e:dc:4c authenticated 00:00:08
Ap_local Away mobility-test/6c:f3:7f:3a:ba:d8/a-HT default tunnel
2001:5000::2677:3ff:fe9e:dc4c 24:77:03:9e:dc:4c authenticated 00:00:08
Ap_local Away mobility-test/6c:f3:7f:3a:ba:d8/a-HT default tunnel
User Entries: 2/2
Curr/Cum Alloc:1/16 Free:1/15 Dyn:2 AllocErr:0 FreeErr:0

(host-HA) #show ip mobile host
Mobile Host List, 1 host(s)
---------------------------
24:77:03:9e:dc:4c
IPv4: 50.50.50.11
IPv6: 2001:5000::2677:3ff:fe9e:dc4c
Roaming Status: Mobile IP Binding (Away), Service time 0 days 00:08:20
Home VLAN 50

(host-HA) #show datapath bridge table 24:77:03:9e:dc:4c
Datapath Bridge Table Entries
-----------------------------
Flags: P - Permanent, D - Deny, R - Roamed Client, M - Mobile, X - Xsec, A - Auth, O -
Outer V
LAN, T - Trusted
MAC VLAN Assigned VLAN QinQ VLAN Destination Flags
----------------- ---- ------------ --------- ----------- -----
24:77:03:9E:DC:4C 4095 50 0 tunnel 9 PMT
24:77:03:9E:DC:4C 50 50 0 tunnel 9 PMTR

(host-HA) #show ip mobile binding
Home Agent Binding list, 1 host(s)
----------------------------------
24:77:03:9e:dc:4c
IPv6: 2001:5000::2677:3ff:fe9e:dc4c
FA Care-of Addr 10.15.44.10, Src Addr 10.15.44.10, HAT HA Addr 10.15.44.60
FA Visiting VLAN 60
Lifetime granted 00:00:40 (40), remaining 00:00:23
Flags T, Registration id D51BA8BC:856865FC
Tunnel id 9, src 10.15.44.60 dest 10.15.44.10, reverse-allowed
```

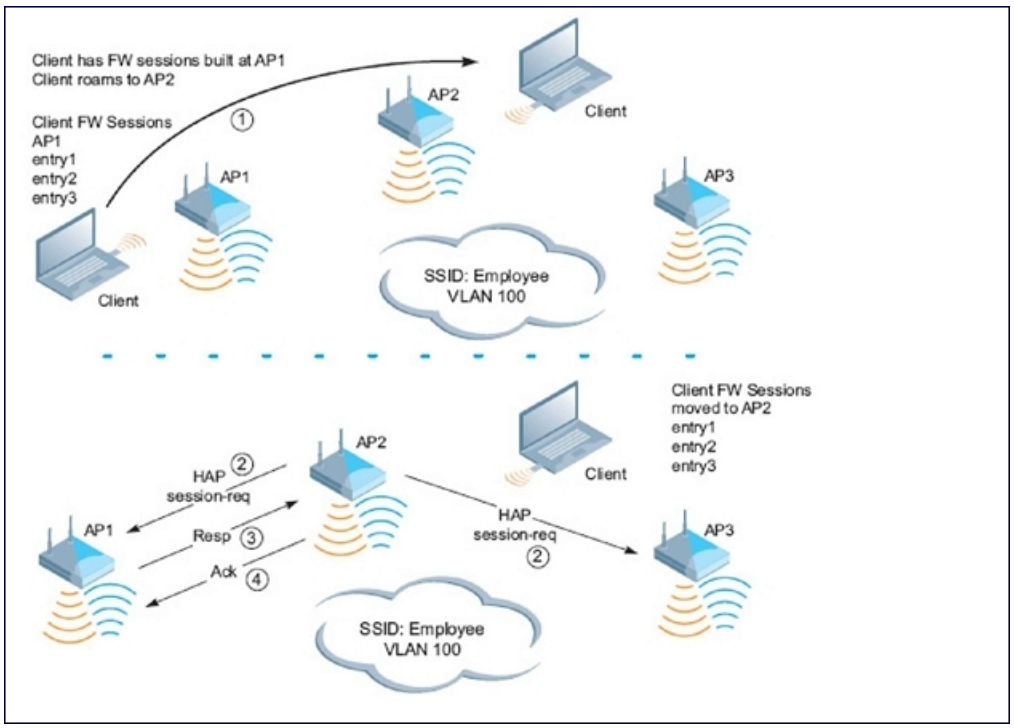
Understanding Bridge Mode Mobility Deployments

In bridge mode deployments, it is possible to deploy more than one AP in a single location. Therefore, APs in bridge forwarding mode support firewall session synchronization, which allows clients to retain their current session and IP address as they roam between different bridge mode APs on the same Layer-2 network.

The bridge mode mobility feature facilitates client mobility on up to 32 Layer-2 connected APs by allowing the APs to communicate and share the user state as the user roams from AP to AP. This

mechanism is always enabled when an AP is set to bridge mode, and it requires that all APs be on the same Layer-2 segment where roaming will occur.

Figure 100 *Bridge Mode Mobility*



The roaming process occurs as follows:

1. A client begins to roam from AP1 and starts an association with AP2.
2. AP2 sends a broadcast message to all APs on the local Layer-2 network, asking if any other AP has a current session state for the roaming client.
3. Only AP1 responds to the broadcast, and sends the current session table of the client.
4. AP2 acknowledges receipt of the session table.
5. AP1 deletes the session state of the client.
6. Roaming is complete.

Monitoring Network Traffic Using IP Flow Information Export

IP Flow Information Export allows clients to easily monitor network traffic to and from the node. This information is cached on the managed device, then exported to an assigned collector server within the node once the table is full or the timer has expired. This information is then logged and stored by the collector server for viewing. Listed below are the tasks to monitor network traffic using IP flow information export:

Enabling IP Flow Information Export

Before enabling IP Flow Information Export, the device must be configured for local management within the node. If the device is not locally managed, the **IPFIX** tab will not be displayed in the WebUI.

The following procedure describes how to enable IP Flow Information Export:

1. In a **Managed Network** node hierarchy, navigate to **Configuration** > **Services** and select the **External Services** tab.
2. Expand **IPFIX** accordion.
3. Click the **Enable IPFIX** toggle switch to enable this setting.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands enable IP Flow Information Export:

```
(host) [md] (config)#ip-flow-export-profile
(host) [md] (ip flow collector profile)#enable
```

Enabling Wireless Export

Starting with ArubaOS 8.0.1.0, IP Flow Information Export supports wireless export. When wireless export is enabled, a new template is defined to gather and export information about WLAN clients, in addition to the standard attributes exported through the existing, pre-defined template.

The wireless attributes include:

- Station MAC
- Station IP
- Station SSID
- AP MAC



If wireless export is enabled, data flows become unidirectional.

The following CLI commands enable wireless export:

```
(host) [mynode] (config) #ip-flow-export-profile
(host) [mynode] (ip flow collector profile) #wireless-export
```

Assigning the Collector Device

When a device belonging to a node exports a cache, it is sent to the designated Collector Device in that node. The Collector Device receives, logs and stores the data from the other devices in the node.

The following procedure describes how to assign the collector IP address:

1. In a **Managed Network** node hierarchy, navigate to **Configuration > Services** and select the **External Services**tab.
2. Expand **IPFIX** accordion.
3. Click the **Enable IPFIX** toggle switch to enable this setting.
4. Enter the IP address of the device in the **Collector IP address** field.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands assign the collector IP address:

```
(host) [md] (config)#ip-flow-export-profile
(host) [md] (ip flow collector profile)#collector-ip <collector ip address>
```

Selecting a Transfer Mode

IP Flow Information Export supports UDP and TCP transfer protocols when sending a cache from a device to the Collector Device.

The following procedure describes how to select a transport mode:

1. In a **Managed Network** node hierarchy, navigate to **Configuration > Services** and select the **External Services**tab.
2. Expand **IPFIX** accordion.
3. Click the **Enable IPFIX** toggle switch to enable this setting.
4. Select a transfer protocol from the **Transport mode** drop-down list.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands select a transport mode:

```
(host) [md] (config)#ip-flow-export-profile
(host) [md] (ip flow collector profile)# transport-protocol<protocol>
```

Assigning a Destination Port

Clients can assign a destination port on the Collector Device to direct incoming data caches from other devices in the node.

The following procedure describes how to ssign a destination port on the Collector Device:

1. In a **Managed Network** node hierarchy, navigate to **Configuration > Services** and select the **External Services**tab.
2. Expand **IPFIX** accordion.
3. Click the **Enable IPFIX** toggle switch to enable this setting.
4. Enter the port number in the **Port** field.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands assign a port on the Collector Device :

```
(host) [md] (config)#ip-flow-export-profile
(host) [md] (ip flow collector profile)#port <port number>
```

Modifying the Flow Cache Size and Interval Settings

The Flow Cache limits when the cache is exported to the Collector Device and can be determined by the size of the cache or the duration of time in the session. When any one of these values is met, the cache is exported and a new one begins.

- **Flow cache size:** The maximum number of entries in a cache before it is exported.
- **Upload interval (all):** The interval (time in minutes) to export active sessions.
- **Upload interval (inactive):** The interval (time in minutes) to export inactive flows.
- **Upload interval template:** The interval (time in minutes) to export templates.

The following procedure describes how to adjust the flow cache size and interval settings:

1. In a **Managed Network** node hierarchy, navigate to **Configuration > Services** and select the **External Services**tab.
2. Expand **IPFIX** accordion.
3. Click the **Enable IPFIX** toggle switch to enable this setting.
4. Enter the maximum number of entries in the **Flow cache size** field.
5. Enter the time interval for an active session in the **Upload interval (all)** field.
6. Enter the time interval for an inactive session in the **Upload interval (inactive)** field.
7. Enter the time interval to export templates in the **Upload interval (template)** field.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands adjusts the Flow cache size and interval export settings:

```
(host) [md] (config)#ip-flow-export-profile
(host) [md](ip flow collector profile)#flow-cache-size<interger>
(host) [md](ip flow collector profile)#upload-all-interval<interger>
(host) [md](ip flow collector profile)#upload-inactive-interval<interger>
(host) [md](ip flow collector profile)#upload-template-interval<interger>
```

Selecting the Observation Domain

The Observation Domain is a value used by the Collector Device to group devices when receiving data sessions.

The following procedure describes how to configure observation domain:

1. In a **Managed Network** node hierarchy, navigate to **Configuration** > **Services** and select the **External Services** tab.
2. Expand **IPFIX** accordion.
3. Click the **Enable IPFIX** toggle switch to enable this setting.
4. Enter the value in the **Observation Domain** field.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure observation domain:

```
(host) [md] (config)#ip-flow-export-profile
(host) [md] (ip flow collector profile)#observation-domain
```

Enabling Mobility Multicast

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 through IGMP messages. Upstream routers use IGMP message information to compute multicast routing tables and determine the outgoing interfaces for each multicast group stream.

When a mobile client moved away from its local network and associated with a VLAN on a foreign managed device (or a foreign VLAN on its own managed device), the client's multicast membership information would not be available at its new destination, and multicast traffic from the client could be

interrupted. However, ArubaOS supports mobility multicast enhancements that provide uninterrupted streaming of multicast traffic, regardless of a client's location. Listed below are the tasks to enable mobility multicast:

Working with Proxy IGMP and Proxy Remote Subscription

The managed device is always aware of the client's location, so the managed device can join multicast group(s) on behalf of that mobile client. This feature, called Proxy IGMP, allows the managed device 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 managed device to maintain a multicast forwarding table). The multicast IGMP traffic originating from the client will instead be sent from the managed device's incoming VLAN interface IP.

The IGMP proxy feature includes both a host implementation and a router implementation. An upstream router sees a managed device running IGMP proxy as a host; a client attached to the managed device sees the managed device as router. When you enable Proxy IGMP, all multicast clients associated with the managed device are hidden from the upstream multicast device or router.



The newer IGMP proxy feature and the older IGMP snooping feature cannot be enabled at the same time, as both features add membership information to multicast group table. For most multicast deployments, you should enable the IGMP Proxy feature on all VLAN interfaces to manage all the multicast membership requirements on the managed device. If IGMP snooping is configured on some interfaces, there is a greater chance that multicast information transfers may be interrupted.

IGMP proxy must be enabled or disabled on each individual interface. To use the IGMP proxy, ensure that the VLANs on the managed device are extended to the upstream router. Enabling IGMP proxy enables IGMP on the interface and sets the querier to the managed device itself. You must identify the managed device port from which the managed device sends proxy join information to the upstream router, and identify the upstream router by upstream port so the managed device can dynamically update the upstream multicast router information.

IGMPv3 Support

ArubaOS supports IGMPv3 functionality that makes Aruba managed devices aware of the Source Specific Multicast and is used to optimize bandwidth of the network. The Source Specific Multicast functionality is an extension of IP multicast where the datagram traffic is forwarded to receivers from only those multicast sources to which the receivers have explicitly joined. By default, the multicast group range of 232.0.0.0 through 232.255.255.255 (232/8) is reserved for Source Specific Multicast by Internet Assigned Numbers Authority.

The IGMPv3 snooping functionality is configured at the edge of the network. The devices that support IGMP snooping listen for the IGMP messages that the host sent to join an IP multicast group. These devices record details of all the hosts and also about the IP multicast group in which a particular host has joined. These devices forward IP multicast traffic to the hosts that have joined the specific multicast group.



The IGMP proxy and IGMP snooping functionalities cannot be enabled on the same VLAN simultaneously.

Configuring Source Specific Multicast Range

The following procedure describes how to configure the Source Specific Multicast range:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration > Interfaces** page and select the **Multicast** tab.
2. In the **IGMP** accordion, enter values for Source Specific Multicast Range in the **SSM range start-ip** and **SSM range mask-ip** fields.
3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

NOTE

The proxy operation will be downgraded to IGMPv2 if any lower version clients are present and reverts to v3 mode if the managed device finds no lower version client joins (reports) for a specified interval of time. In the downgraded proxy operation the Source Specific Multicast semantics is not applicable for the particular VLAN.

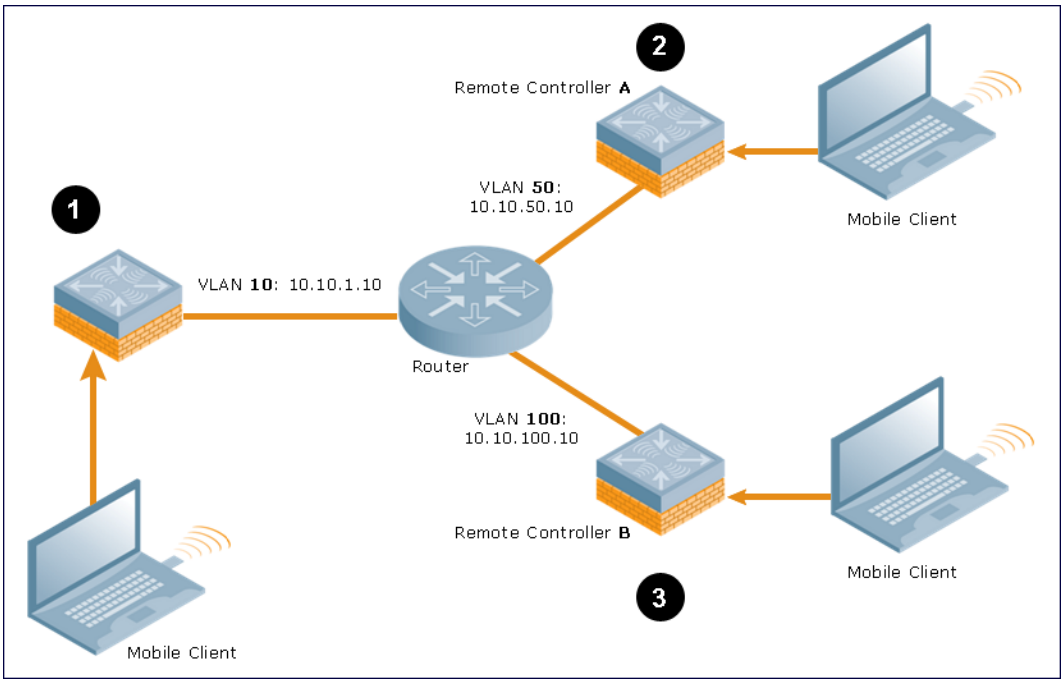
The following CLI commands configure the Source Specific Multicast range:

```
(host) [md] (config) # ip igmp
(host) [md] (config-ip)# ssm-range <startip> <maskip>
```

Working with Inter managed device Mobility

When a client moves from one managed device to another, multicast traffic migrates as follows:

Figure 101 *Inter-managed device Mobility*



1. The managed device uses its VLAN 10 IP address to join multicast group 1 on behalf of a mobile client.

2. The mobile client leaves its managed device and roams to VLAN 50 remote managed device A.

Remote managed device A locates the mobile client's managed device and learns about the client's multicast groups. Remote managed device A then joins group 1 on behalf the mobile client, using its VLAN 50 source IP. Upstream multicast traffic from the roaming client is sent to the managed device over a Layer-2 GRE tunnel. The remote managed device will receive downstream multicast traffic and send it to the mobile client.



The Layer-2 GRE Tunnel implementation of the IP mobility functionality is supported only on ArubaOS versions 6.2.0.0 or later, and is not backward compatible with the earlier implementation. ArubaOS supports only v4 mobility and does not support IPv6 Layer-3 mobility.

Meanwhile, the managed device checks to see if other local clients require group 1 traffic. If no other clients are interested in group 1, then the managed device will leave that group. If there are other clients using that group, the managed device will continue its group 1 membership.

3. Now the mobile client leaves remote managed device A and roams to VLAN 100 on remote managed device B. Remote managed device B locates the mobile client's managed device and learns about the client's multicast groups. Remote managed device B then joins group 1 on behalf the roaming mobile client 1, using its VLAN 100 IP address.

Both the managed device and remote managed device A will verify if any of their other clients require group 1 traffic. If none of their other clients require group 1, then that managed device will leave the group. (If the managed device leaves the group, it will also notify remote managed device A). If either managed device has other clients using that group, that managed device will continue its group 1 membership.

Configuring Mobility Multicast

The following procedure describes how to enable IGMP or IGMP snooping on this interface, or configure a VLAN interface for uninterrupted streaming of multicast traffic:

1. In a **Managed Network** node hierarchy, navigate to the **Configuration** > **Interfaces** page and select the **VLANs** tab.
2. Select the VLAN name from the **VLANs** table
3. In the **VLANs > <VLAN name>** table, select a VLAN ID you want to configure mobility multicast.
4. Select **IPv4** tab and expand **IGMP** accordion.
5. Select **snooping** from the **IGMP** drop-down list to save bandwidth and limit the sending of multicast frames to only those nodes that need to receive them.
6. Select **proxy** from the **IGMP** drop-down list to enable the router to discover the presence of multicast listeners on directly-attached links.
7. In the **Proxy Interface** field, select the **Interface** or **Port Channel** option and select the value from the drop-down list.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure mobility multicast:

```
(host) [md] (config)#interface vlan <vlan>
(host) [md] (config-submode)#ip igmp proxy [{port-channel|gigabitethernet}
<slot/module/port>]|[snooping]
```

Multicast Group Limit

The following table describes the maximum multicast group limit per managed device platform.



Maximum multicast group is the sum of IPv4 IGMP and IPv6 MLD groups.

Table 152: *Multicast Group Limits*

Platform	Multicast Group Limit
7005	128
7010	256
7024	256
7030	512
7200 Series	4096

Chapter 31
External Firewall Configuration

In many deployment scenarios, an external firewall is situated between Aruba devices. This chapter describes the network ports that need to be configured on the external firewall to allow proper operation of the Aruba network. You can also use this information to configure session ACLs to apply to physical ports on the managed device for enhanced security. However, this chapter does not describe requirements for allowing specific types of user traffic on the network.



A managed device 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 managed device.

Topics in this chapter include:

- Understanding Firewall Port Configuration in Aruba Devices on page 831
- Enabling Network Access on page 832
- Ports Used for VIA on page 832
- Configuring Ports to Allow Other Traffic Types on page 833

Understanding Firewall Port Configuration in Aruba Devices

This section describes the network ports that need to be configured on the firewall to allow proper operation of the network.

Communication Between Managed Devices

Configure the following ports to enable communication between any two managed devices:

- IPsec (UDP port 4500) for communication between Mobility Master and a managed device.
- IPsec (UDP ports 500 and 4500) and ESP (protocol 50). PAPI between Mobility Master and a managed device 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 managed device
- IKE (UDP 500)
- ESP (protocol 50)
- NAT-T (UDP 4500)

Communication Between APs and the Managed Device

APs use Trivial File Transfer Protocol (TFTP) during their initial boot to grab their software image and configuration from the managed device. After the initial boot, the APs use FTP to retrieve their software images and configurations from the managed device. In many deployment scenarios, an external firewall is situated between various Aruba devices.

Configure the following ports to enable communication between an AP and the managed device:

- PAPI (UDP port 8211). If the AP uses DNS to discover the LMS managed device, the AP first attempts to connect to the managed device. (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 managed device.
- FTP (TCP port 21)
- TFTP (UDP port 69). All campus APs; If there is no local image on the AP or if the image needs to be upgraded (for example, a new AP), the AP will use TFTP to retrieve the initial image. For remote APs, upgrade the image only by FTP and not TFTP.
- SYSLOG (UDP port 514)
- PAPI (UDP port 8211)
- GRE (protocol 47)
- Control Plane Security (CPsec) uses UDP port 4500

Communication Between Remote APs and the Managed Device

Configure the following ports to enable communication between a Remote AP (IPsec) and a managed device:

- NAT-T (UDP port 4500)
- TFTP (UDP port 69)



TFTP is not needed for normal operation. If the remote AP loses its local image for any reason, it will use TFTP to download the latest image.



In an IPv6 deployment, you must configure UDP port 500 on the firewall to establish an IPv6 IPsec connection for Remote APs.

Enabling Network Access

This section describes the network ports that need to be configured on the firewall to manage the Aruba network.

For WebUI access between the network administrator's computer (running a web browser) and a managed device:

- HTTP (TCP ports 80 and 8888) or HTTPS (TCP ports 443 and 4343).
- SSH (TCP port 22) or TELNET (TCP port 23).

Ports Used for VIA

The following ports are used with Aruba VIA.

- For the reachability/trusted network check, use port 443.
- For the IPsec connection, use port 4500.
- To allow ISAKMP, use port 500.
- To enable NAT-T, use port 4500.

Configuring Ports to Allow Other Traffic Types

This section describes the network ports that need to be configured on the firewall to allow other types of traffic in the Aruba network. You should only allow traffic as needed from these ports.

- For logging: SYSLOG (UDP port 514) between the managed device and syslog servers.
- For software upgrade or retrieving system logs: TFTP (UDP port 69) or FTP (TCP ports 21 and 22) between the managed device and a software distribution server.
- If the managed device is a PPTP VPN server, allow PPTP (UDP port 1723) and GRE (protocol 47) to the managed device.
- If the managed device is an L2TP VPN server, allow NAT-T (UDP port 4500), ISAKMP (UDP port 500) and ESP (protocol 50) to the managed device.
- If a third-party NMS is used, allow SNMP (UDP ports 161 and 162) between the NMS and all managed devices.
- For authentication with a RADIUS server: RADIUS (typically, UDP ports 1812 and 1813, or 1645 and 1646) between the managed device and the RADIUS server.
- For authentication with an LDAP server: LDAP (UDP port 389) or LDAPS (UDP port 636) between the managed device and the LDAP server.
- For authentication with a Radsec Server—RADIUS over TLS): TCP port 443 between the managed device and the RADIUS or Radsec Proxy server.
- For authentication with a TACACS+ server: TACACS (TCP port 49) between the managed device and the TACACS+ server.
- For NTP clock setting: NTP (UDP port 123) between all managed devices and NTP 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 managed device configuration.
- For External Services Interface (ESI): ICMP (protocol 1) and syslog (UDP port 514) between a managed device and any ESI servers.
- For XML API: HTTP (TCP port 80) or HTTPS (TCP port 443) between a managed device and an XML-API client.

Starting from ArubaOS 8.0.1, a minor security enhancement is made to Process Application Programming Interface (PAPI) messages. With this enhancement, PAPI endpoints authenticate the sender by performing a sanity check of the incoming messages using MD5 (hash).

All PAPI endpoints—access points, Mobility Access Switches, controllers, Analytics and Location Engine (ALE), Aruba Switches, HPE-ArubaOS Switch-based switches, AirWave, Mobility Master and Managed Devices—must use the same secret key.



The same PAPI key must be configured for the Mobility Master and the managed device.

The PAPI Enhanced Security configuration provides protection to Aruba devices, AirWave, and ALE against malicious users sending fake messages that results in security challenges.



PAPI Enhanced Security does not solve all the PAPI security issues.

Topics in this chapter include:

- Interoperability
- Configuring PAPI Enhanced Security
- Verifying PAPI Enhanced Security

Interoperability

The following list of references provides the Aruba devices interoperability information with respect to PAPI Enhanced security feature:

- For information on interoperability with AirWave, refer to the *AirWave 8.2.0.3 Release Notes*.
- For information on interoperability with Analytics and Location Engine (ALE), refer to the *Analytics and Location Engine 2.0.0.6 Release Notes*.
- For interoperability with Mobility Access Switches, refer to the *ArubaOS 7.4.1.5 Release Notes*.
- For interoperability with HPE-ArubaOSSwitch-based switches, refer to HP's *Management Configuration Guide 16.02*.

AirWave Management Platforms–AMP 8.0.11.2 and AMP 8.2.3–support PAPI Enhanced Security.

Configuring PAPI Enhanced Security

By default, the PAPI Enhanced Security configuration is disabled. If there is no configured key, the default key is used for authentication.

The following CLI commands configure PAPI Enhanced Security:

```
(host)[mynode] (config) #papi-security
(host)[mynode] (PAPI Security Profile) #?
enhanced-security    Enable or disable the use of enhanced security mode
key                  Key used to authenticate messages between systems.
Length must be between 10 and 64 characters. Use 'no key' to revert to the default key.
```

```
no                    Delete Command
(host)[mynode] (PAPI Security Profile) #enhanced-security
(host)[mynode] (PAPI Security Profile) #write memory
Saving Configuration...
Partial configuration for /mm/mynode
-------------------------------------
Contents of : /flash/config/partial/53/p=sc.cfg
papi-security
enhanced-security
!
Configuration Saved.
```

If PAPI Enhanced Security is enabled with Layer 2 redundancy and VRRP, a static route needs to be added on managed devices to the master IP (virtual IP) address. The static route needs to be added via its conductor redundancy ipsec-map using the command `ip route <destip> <mask> ipsec <map-name>`.

Verifying PAPI Enhanced Security

The following CLI command verifies the status of the PAPI Enhanced Security configuration:

```
(host)[node] (config) #show papi-security
PAPI Security Profile
---------------------
Parameter               Value
---------               -----
PAPI Key                ********
Enhanced security mode Disabled
```

To view the statistics of transmitted, received, and denied messages, three additional output parameters are introduced in the **show ipc statistics** command output.

- Tx Sign—the number of messages which were signed before transmitting
- Rx Sign—the number of messages validated through digest validation
- Rx Denied—the number of messages denied due to incorrect digest

```
(host) [mynode] #show ipc statistics app-name sapm
Local Statistics
To application     Tx Msg   Tx Blk  Tx Ret Tx Fail  Rx Ack Rx Msg Rx Drop
Layer2/3                4       0       0       0       0      2    0
Multicast DNS Lis       0       0       0       0       0      3    0
License Manager         2       2       0       0       2      2    0
Profile Manager         1       0       0       0       1      1    0
NEW_CLI_START           2       0       0       0       2      3    0
Authentication          0       0       0       0       0      1    0
Syslog Manager          4       4       0       0       4      0    0
Configuration Man       3       0       0       0       0     19    0
     Rx Err   Tx Ack   Tx Sign   Rx Sign  Rx Denied Rx Silent Drops
         0        0         0         0         0             0
         0        0         0         0         0             0
         0        0         0         0         0             0
         0        0         0         0         0             0
         0        0         0         0         0             0
         0        0         0         0         0             0
         0        0         0         0         0             0
         0        0         0         0         0             0
Kernel PAPI Statistics
RxSockbufSize RxSockbufHimark CurRxQLen MaxRxQLen Drops
16777216                1152         0         1         0
Remote Device 10.4.176.95 Statistics
To application      Tx Msg   Tx Blk   Tx Ret  Tx Fail   Rx Ack   Rx Msg
```

```
SAPM                    2565         0         0         0         0     2667
Rx Drop    Rx Err   Tx Ack   Tx Sign   Rx Sign   Rx Denied   Rx Silent Drops
      0         0        0         0         0           0               0

Remote Device 172.200.13.3 Statistics
To application       Tx Msg    Tx Blk    Tx Ret   Tx Fail    Rx Ack    Rx Msg
SAPM                   2569         0         0         0         0     2569
Rx Drop    Rx Err    Tx Ack   Tx Sign   Rx Sign   Rx Denied   Rx Silent Drops
      0         0         0         0         0           0               0
Allocated Buffers    4
Static Buffers       0
Static Buffer Size   1476
```

The User-Identification (User-ID) feature of the Palo Alto Networks firewall allows network administrator to configure and enforce firewall policies based on users and user groups. The User-ID identifies the user on the network based on the IP address of the device to which the user is logged in. Additionally, a firewall policy can be applied based on the type of device the user is using to connect to the network. Since the Mobility Master maintains the network and user information of clients in the network, it is the best source to provide information for the User-ID feature of the PAN firewall.



The procedures in this chapter describe the steps to integrate a Palo Alto Networks firewall with a Mobility Master or managed device. **For additional details on configuring PAN firewall integration, see**Managed Device Feature Overview on page 226.

This feature supports the following interactions with Palo Alto Networks firewall servers running PAN-OS 5.0 or later:

- Send login events for the client to the PAN firewall with its IP address, username, and device type, when classified.
- Send logout events for the client to PAN firewalls with its IP address.

The following must be configured on the PAN Firewall:

- An admin account must be created on the PAN firewall to allow the managed device to send data to the PAN firewall. This account must match the account added in the PAN profile on the managed device. The built-in admin account can be used for this purpose, but that is not recommended. It is better to create a new admin account used solely for the purpose of communications between the managed device and PAN firewall.
- Pre-configuration of PAN Host Information Profile objects and HIP-profiles on the PAN Firewall to support a device-type based policy.
- To enable these features, the following must be configured on the managed device:
  - The system-wide PAN profile must be properly configured and made active on the managed device.
  - The **pan-integration** parameter in the AAA profile to which the client is associated must be enabled.
  - For VPN clients, enable the **pan-integration** parameter in the VPN authentication profile to which the client is associated.
  - For VIA clients, enable the **pan-integration** parameter in the VIA authentication profile to which the client is associated.



PAN Firewall Integration does not support bridge forwarding mode.

Pre-configuration on the PAN Firewall

Before PAN Firewall configuration can be completed on the managed device, some configurations must be completed on the PAN Firewall.

Certificate Management

The issuer certificate of the x509 server certificate used by the PAN firewall must be imported by Mobility Master as a trusted CA in order to establish a secure HTTPS connection between the firewall and the managed device.

User-ID Support

The administrator must configure firewall policies based on the username and/or user group. Additionally, correct configuration of the connection to directory servers is required for user group based policies on the PAN firewall.

Device-Type Based Policy Support

Managed devices support a limited number of device types. The identified device type associated with each IP user is sent to the PAN through the **client-version** field, with the **host-info** category of the HIP report. PAN administrators must create these HIP objects, which filter the HIP reports sent from the managed device to support device-based firewall policies.

Table 153 lists the HIP objects with a specified **Is Value** in the **Client Version** field, which must be preconfigured on the PAN firewall.

Table 153: *HIP Objects*

Client Version Is Value
Android
Apple
AppleTV
BlackBerry
Chrome OS
iPad
iPhone
iPod
Kindle
Linux
Nintendo
Nintendo 3DS
Nintento Wii
Nook

Client Version Is Value
OS X
PlayStation
PS Vita
PS3
PSP
RIM Tablet
Roku
Symbian
webOS
Win 7
Win 8
Win 95
Win 98
Win 2000
Win CE
Win ME
Win NT
Win Server
Win Vista
Win XP
Windows
Windows Mobile
Windows Phone 7

Configuring PAN Firewall Integration

A PAN profile must be created on the managed device. Multiple PAN profiles can be configured and saved on the managed device, but only one profile can be active at a time.



The following procedures describe the steps to integrate a Palo Alto Networks firewall policy using Mobility Master. For additional details on configuring PAN firewall integration, see Managed Device Feature Overview on page 226

Listed below are the tasks to be executed to for PAN Firewall integration:

Creating PAN Profiles

The first step in configuring PAN firewall integration is to create PAN profiles. This profile provides the managed device with the information required for connecting to and interacting with the specified PAN firewall.



This configuration is only performed and available on the Mobility Master. The configuration is pushed to all connected to the managed devices.

The following procedure describe how to create a new PAN profile and add PAN firewalls:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System** page.
2. Select the **Profiles** tab.
3. Expand **Other Profiles** in the **All Profiles** list, and then select **Palo Alto Network Servers**.
4. Click **+**.
5. Type the profile name into the **Profile Name** field, then click **+**.
6. In the **Host** field, enter the host IP address or the host name of the PAN firewall
7. In the **Portnum** field, enter the port number (1 – 65535) of the PAN firewall.
8. In the **Username** field, enter the admin username or the PAN firewall. The username must be between 1 and 255 bytes in length. The username must match the Admin account previously created on the PAN firewall.
9. In the **Passwd** field, enter the password of the username in PAN firewall. The password is between 6 and 100 bytes in length. The password must match the Admin account previously created on the PAN firewall.
10. In the **Retype** field, re-enter the password entered in the previous step.
11. Click **OK**.
12. Click **Submit**.
13. Select **Pending Changes**.
14. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Adding PAN Firewalls

The following procedure describe how to add additional PAN firewalls to an existing PAN profile:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Systems** page.
2. Select the **Profiles** tab.
3. Expand **Other Profiles** in the **All Profiles** list, and then select the PAN profile under **Palo Alto Networks Servers**.
4. Click **+** to add a PAN firewall to the PAN profile.
5. Enter the **Host** (IP address or hostname) of the PAN firewall
6. Enter the **Portnum** (port number between 1–65535) of the PAN firewall.
7. Enter the **Username** of the PAN firewall. The username must be between 1 and 255 bytes in length. The username must match the admin account previously created on the PAN firewall.
8. Enter the **Passwd** of the username in PAN firewall. The password must be between 6 and 100 bytes in length. The password must match the admin account previously created on the PAN firewall.
9. In the **Retype** field, re-enter the password.
10. Click **OK**.
11. Click **Submit**.

12. Select **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI snippet creates a new PAN profile:

```
(host)[mynode](config) #pan profile <profile-name>
firewall host <host> port <port> username <username> passwd <password>
```

Activating a PAN Profile

The following procedure describe how to activate a PAN Firewall profile:



This configuration must be completed on each managed device.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System** page
2. Select the **Profiles** tab.
3. Expand **Other Profiles** in the **All Profiles** list, and then select the **Palo Alto Networks Active >Active Palo Alto Networks** profile.
4. Select the profile from the **Palo Alto Networks Servers Profile** drop-down list, and then click **+.**
5. Enter the **Host** (IP address or hostname) of the PAN firewall
6. Enter the **Portnum** (port number between 1 – 65535) of the PAN firewall.
7. Enter the **Username** of the PAN firewall. The username must be between 1 and 255 bytes in length. The username must match the admin account previously created on the PAN firewall.
8. Enter the **Passwd** for the PAN firewall. The password must be between 6 and 100 bytes in length. The password must match the admin account previously created on the PAN firewall.
9. Re-enter the password.
10. Click **OK**.
11. Click **Submit**.
12. Select **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command activates a PAN Firewall profile:

```
(host) [mynode] (config) #pan active-profile
profile <profile-name>
```

Configuring PAN Firewall Portal

The following procedure describe how to configure the PAN firewall portal:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System** page.
2. Select the **Profiles** tab.
3. Expand **Other Profiles** in the **All Profiles** list, and then select the **Configure Palo Alto Network options** profile.
4. Enter the **host** (IP address or hostname) of the PAN firewall
5. Enter the name of the trusted CA certificate under **cacert_name**.
6. Enter the username of the PAN firewall in the **uname** field. The username must be between 1 and 255 bytes in length. The username must match the admin account previously created on the PAN firewall.

7. Enter the **passwd** for the PAN firewall. The password must be between 6 and 100 bytes in length. The password must match the admin account previously created on the PAN firewall.
8. Re-enter the password.
9. Click **Submit**.
10. Select **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Enabling PAN Firewall Integration

The following procedure describe how to enable PAN firewall integration on the AAA profile to which the client is associated:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System** page.
2. Select the **Profiles** tab.
3. Expand **Wireless LAN** in the **All Profiles** list, and then select **AAA**.
4. Select an **AAA** profile.
5. Select the **PAN Firewall Integration** check box.
6. Click **Submit**.
7. Select **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command enables PAN firewall integration on the AAA profile:

```
(host) [mynode] (config) #aaa profile <aaa profile-name>
pan-integration
```

Enabling PAN Firewall Integration for VIA Clients

For VIA clients, PAN firewall integration must be enabled on the VIA authentication profile that is associated with the client.

The following procedure describe how to enable PAN firewall integration for VIA clients:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System** page.
2. Select the **Profiles** tab.
3. Expand **Other Profiles** in the **All Profiles** list, and then select **VIA Authentication**.
4. Select a **VIA Authentication** profile.
5. Select the **PAN Firewalls Integration** check box.
6. Click **Submit**.
7. Select **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command enables PAN firewall integration for VIA clients:

```
(host) [mynode] (config) #aaa authentication via auth-profile <profile-name>
pan-integration
```

Enabling PAN Firewall Integration for VPN Clients

For VPN clients, PAN firewall integration must be enabled on the VPN authentication profile that the client is associated with.

The following procedure describe how to enable PAN firewall integration for VPN clients:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System** page.
2. Select the **Profiles** tab.
3. Expand **Wireless LAN** in the **All Profiles** list, and then select **VPN Authentication**.
4. Select a **VPN Authentication** profile.
5. Select the **PAN Firewalls Integration** check box.
6. Click **Submit**.
7. Select **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command enables PAN firewall integration for VPN clients:

```
(host) [mynode] (config) #aaa authentication vpn default
pan-integration
```

Related Commands

Use the following CLI commands to view details for your PAN firewall configuration:

```
(host) [mynode] #show pan activate-profile
(host) [mynode] #show pan debug [uid-table slot-num <slot-num> [starting-rec
<starting-rec>]]
(host) [mynode] #show pan-gp gateway-info
(host) [mynode] #show pan-gp portal-info
(host) [mynode] #show pan-options
(host) [mynode] #show pan profile
(host) [mynode] #show pan profile <profile-name>
(host) [mynode] #show pan state
(host) [mynode] #show pan statistics
(host) [mynode] #show profile-list pan profile [start <start>] [page <page>]
(host) [mynode] #show references pan active-profile [start <start>] [page <page>]
(host) [mynode] #show references pan profile <profile-name> [start <start>] [page
<page>]
```

Chapter 34
Remote Access Points

The Secure Remote Access Point Service allows AP users, at remote locations, to connect to an Aruba Managed Device over the Internet. As the Internet is involved, data traffic between the Managed Device and the remote AP is VPN encapsulated. That is, the traffic between the Managed Device and AP is encrypted. Remote AP operations are supported on all of Aruba's APs.

Topics in this chapter include:

- About Remote Access Points on page 844
- Configuring the Secure Remote Access Point Service on page 845
- Deploying a Branch or Home Office Solution on page 852
- Bringing up Certificate-Based Remote AP in VMC on page 857
- Enabling Remote AP Advanced Configuration Options on page 859
- Understanding Split Tunneling on page 873
- Understanding Bridge on page 881
- Provisioning Wi-Fi Multimedia on page 886
- Reserving Uplink Bandwidth on page 886
- Provisioning 4G USB Modems on Remote Access Points on page 887
- Converting an Instant AP to Remote AP or Campus AP on page 892
- Enabling Bandwidth Contract Support for Remote APs on page 893

About Remote Access Points

Remote APs connect to a managed device using XAuth or IPsec. AP control and 802.11 data traffic are carried through this tunnel. Secure Remote AP Service extends the corporate office to the remote site. Remote users can use the same features as corporate office users. For example, VoIP applications can be extended to remote sites while the servers and the PBX remain secure in the corporate office.

For bothRemote APs and Campus APs, tunneled SSIDs will be brought down eight seconds after the AP detects that there is no connectivity to the Managed Device. However, Remote AP bridge-mode SSIDs are configurable to stay up indefinitely (always-on or persistent). For Campus AP bridge-mode SSIDs, the Campus AP will be brought down after the keepalive times out (default 3.5 minutes).

Secure Remote AP Service can also be used to secure control traffic between an AP and the managed device in a corporate environment. In this case, both the AP and managed device 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 managed device reside in a private network which secures AP-to-Managed Device communication. (This deployment is recommended when AP-to-Managed Device communications on a private network need to be secured.) In this scenario, the Remote AP uses the managed device's IP address on the private network to establish the IPsec VPN

tunnel.

- Deployment Scenario 2: The Remote AP is on the public network or behind a NAT device and the managed device 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 managed device in the DMZ. The Remote AP uses the managed device's IP address on the public network to establish the IPsec VPN tunnel.
- Deployment Scenario 3: The Remote AP is on the public network or behind a NAT device and the managed device is also behind a NAT device. (This deployment is recommended for remote access.) The Remote AP must be configured with the tunnel termination point, which must be a publicly-routable IP address. In this scenario, the Remote AP uses the public IP address of the corporate firewall. The firewall forwards traffic to an existing interface on the managed device (The firewall must be configured to pass NAT-T traffic (UDP port 4500) to the managed device).

In any of the described deployment scenarios, the IPsec VPN tunnel can be terminated on a managed device, with a managed device located elsewhere in the corporate network. The remote AP must be able to communicate with the managed device after the IPsec tunnel is established. Make sure that the L2TP IP pool configured on the managed device (from which the Remote AP obtains its address) is reachable in the managed device network by the managed device.

It is not recommended to place a Remote AP in the same subnet as its terminating controller. Each Remote AP is deployed at a remote location that is connected over a multi-hop public or private IP network where a direct Layer 2 path to the Mobility Controllers in the data center is not possible. As a best practice, always place an IP router between the APs and the Mobility Controllers as it establishes Layer 2 fault domains.

Configuring the Secure Remote Access Point Service

The tasks for configuring an Aruba Access Point as a Secure Remote Access Point Service are:

- Configure a public IP address for the Managed Device.

  You must install one or more AP licenses in the Managed Device. 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 Managed Device.
- Configure the VPN server on the Managed Device. The remote AP will be a VPN client to the server.
- Provision the AP with IPsec settings, including the username and password for the AP, before you install it at the remote location. You can also provision the Remote AP using the ZTP method. For more information, see Provisioning 4G USB Modems on Remote Access Points on page 887.

Configuring a Public IP Address for the Managed Device

The remote AP requires an IP address to which it can connect to establish a VPN tunnel to the Managed Device. This can be either a routable IP address you configure on the Managed Device, or the address of an external router or firewall that forwards traffic to the Managed Device. The following procedure describes how to create a DMZ address on the Managed Device.

The following procedure describes how to configure a public IP address for the managed device:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Interfaces** > **VLANs** page.
2. Click **+** to add a VLAN.
3. Enter the **VLAN name** and **VLAN ID/Range**.

4. Click **Submit**.
5. Click the Vlan Id created. The **VLANs>name** table is displayed.
6. Click the Vlan Id from **VLANs>name** table.
7. Click **Edit** on **Port Members**.
8. Click > to select the port that belongs to this VLAN .
9. Click **OK**.
10. Click **Submit**.
11. Click **IPv4** tab.
12. Enter the IPv4 address in the **IPv4 address** field.
13. Click **Submit**.
14. Click **Pending Changes**.
15. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure a public IP address for the managed device:

```
(host) [md] (config) #vlan <id>
(host) [md] (config) #interface vlan <id>
(host) [md] (config-submode) #ip address <ipaddr> <ipmask>
```

Configuring the NAT Device

Communication between the AP and the secure Managed Device uses the UDP 4500 port. When both the Managed Device 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 Managed Device to ensure that the remote AP boots successfully.

Configuring the VPN Server

This section describes how to configure the IPsec VPN server on the Managed Device. For more details, see Virtual Private Networks on page 459. The remote AP will be a VPN client that connects to the VPN server on the Managed Device.

1. The following procedure describes how to configure the VPN server:
2. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Services** > **VPN** tab.
3. Click **IKEv1** accordion.
4. Click **L2tp** check box to enable L2tp.
5. Select **PAP (Password Authentication Protocol)** check box for **Auth protocols**.
6. To configure the L2TP IP pool, click **General VPN** option.
7. Click **+** in the **Address Pools** table.
8. Enter the **Pool name** to configure the L2TP pool from which the APs will be assigned addresses.
9. Enter the value of the **Start address (ipv4/v6)** and **End address (ipv4/v6)** fields.
10. Click **Submit**.



The size of the pool should correspond to the maximum number of APs that the Mobility Master is licensed to manage.

11. To configure an ISAKMP encrypted subnet and PSK, click the **Shared Secrets** accordion.
12. Click **+** in the **IKE Shared Secrets** table.

13. In the **Create IKE Group** table, enter the value for **Shared key** and re-enter the key in **Retype shared key**.
14. Click **Submit**.
15. Click **Pending Changes**.
16. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the VPN server:

```
(host) [md] (config) #vpdn group l2tp
(host) [md] (config-submode) #ppp authentication PAP
(host) [md] (config-submode) #ip local pool <pool_name> <pool_start_address> <pool_end_address>
(host) [md] (config) #crypto isakmp key <keystring> address <ipaddr> netmask <mask>
```

Configuring CORP DNS Server

EST servers are generally deployed in a corporate network and EST enrollment can take place only if the DNS requests are resolved. Starting from ArubaOS 8.5.0.0, Remote Access Points will use CORP DNS server for EST enrollment. This will enable the Remote AP to use its own DNS server and not the one provided by the Internet Service Provider, because the ISP provided DNS server will not be reachable within a corporate network and hence will result in failure of DNS requests. EST enrollment cannot not take place if the DNS requests fail. Thus it is necessary for Remote AP to configure corp DNS server for successful EST enrollment.



CORP DNS server can be configured only for Remote Access Points.

This feature supports both IPv4 and IPv6 addresses and only two CORP DNS servers can be configured.

The following procedure describes how to configure CORP DNS server:

1. In the **Managed Network** node of the hierarchy, navigate to **Configuration** > **System > Profiles.**
2. Expand **AP**. Select **AP system profile**. Click **+** to create a new profile.
3. Enter a **Profile Name.**
4. Expand the **Remote AP** accordion.
5. To configure an IPv4 address, click **+** in **Remote-AP CORP DNS server** and enter the **IP address** in the **Remote-AP CORP DNS server** text box. Click **OK.**
6. To configure an IPv6 address, click **+** in **Remote-AP CORP DNS server IPV6** and enter the **IP address** in the **Remote-AP CORP DNS server IPV6** text box. Click **OK.**
7. Click **Submit**.
8. Select **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure a CORP DNS server:

```
(host) [mynode] (config) #ap system profile <profile-name>
(host) [mynode] (AP system profile <profile-name>) #rap-corp-dns-server <ipv4 address>
(host) [mynode] (AP system profile <profile-name>) #rap-corp-dns-server_ ipv6  <ipv6
address>
```

CHAP Authentication Support over PPPoE

RAPs can now establish a PPPoE session with a PPPoE server at the ISP side and get authenticated using the CHAP. The PPPoE client running on a Remote AP is capable of handling the CHAP authentication requests from the PPPoE server.

NOTE

The PPPoE client selects either the PAP or the CHAP credentials for the Remote AP authentication depending upon the request from the PPPoE server.

The following procedure describes how to configure CHAP:

1. In the **Managed Network** node of the hierarchy, navigate to the **Configuration** > **Access Points** > **Remote APs** tab. The list of discovered APs are displayed on this page.
2. Select the AP you want to configure using CHAP and click **Provision**.
3. In the popup window click **Continue and Reboot**.
4. Click **Uplink** tab and enter the **CHAP Secret** .

NOTE

You can use all the special characters except question mark (?) and the space can be used within double quotes (“ ”).

5. Enter the CHAP Secret again in the **Retype** text box for confirmation.
6. Click **Submit** and **Reboot.**
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configures CHAP:

```
(host) [md] (config) #provision-ap pppoe-chap-secret <KEY>
(host) [md] (config-submode) #reprovision ap-name <name>
```

Configuring Certificate Remote AP

You can configure the remote AP to use the internal certificate for authentication.

The following procedure describes how to configure the certificate Remote AP:

1. In the **Managed Network** node of the hierarchy, navigate to **Configuration** > **Access Points > Remote APs** tab.
2. Select a check box next to the AP Name in the Remote AP table click **Provision**.
3. In the **General** tab, select **Certificate** from the **Authentication methods** drop-down list.
4. Click **Submit** to apply the configuration and reboot the AP as certificate Remote AP.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures the certificate Remote AP:

```
(host) [mynode] (config) #whitelist-db rap add
(host) [mynode] (config) #whitelist-db rap add mac-address <mac-address>
```

Creating a Remote AP Whitelist

If you use the ZTP method to provision the certificate Remote AP, then you must create a Remote AP whitelist. For more information on ZTP of the Remote AP, see Provisioning 4G USB Modems on Remote Access Points on page 887.

Remote AP whitelist is the list of approved APs that can be provisioned on your Managed Device.

The following procedure describes how to create a Remote AP whitelist:

1. In the **Managed Network** node of the hierarchy, navigate to **Configuration** > **Access Points > Whitelist** tab.
2. Click **Remote AP Whitelist** tab.
3. Click **+** and provide the following details:
   - **MAC Address** - Enter the MAC address of the AP.
   - **AP Group** - Select a group to add the AP.
   - **AP Name** - Enter a name for the AP. If you do not enter an AP name, the MAC address will be used instead.
   - **Description** - Enter a text description for the AP.
4. Click **Submit** to add the remote AP to the whitelist.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Configuring PSK Remote AP

You can use PSK authentication to provision an individual remote AP or a group of remote APs using an IKE PSK.

The following procedure describes how to configure PSK authentication for Remote AP:

1. In the **Managed Network** node hierarchy, navigate to **Configuration** > **Access Points > Remote APs** tab.
2. Select the required Remote AP and click **Provision**.
3. In the **General** tab, select **Pre-shared Key** from the **Authentication methods** drop-down list.
4. Enter and confirm the IKE PSK.
5. Select **Global User Name/password** or a **Per AP User Name/Password** from**User credential assignment** drop-down list.
   a. If you use the **Per AP User Names/Passwords** option, each Remote AP is given its own username and password.
   b. If you use the **Global User Name/Password** option, all selected RAPs are given the same (shared) username and password.
6. Enter the user name, and enter and confirm the password. If you want the managed device to automatically generate a user name and password, select **Use Automatic Generation**. If this option is not selected, the user has to enter it manually.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Add the User to the Internal Database

The following procedure describes how to add the user to the Internal database:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Authentication** > **Auth Servers** tab.
2. From the **All Servers** table, select **Internal** under the **Name** field.
3. Click **+** in the **Users** tab. The **Internal Server > Add New User** page is displayed.
4. Enter the **User Name** and **Password**.
5. Click **Enabled** check box to activate this entry on creation.

6. Click **Submit**. Note that the configuration does not take effect until you perform this step.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands add the user to the Internal database:

Execute the following command:

```
(host) [mynode] (config) #local-userdb add username rapuser1 password <password>
```

Remote AP Static Inner IP Address

The Remote AP static inner IP address feature assigns a static inner IP address to a Remote AP. A new *remote-IP address* parameter is added to the existing configuration commands.



You cannot configure the IP Address parameter using the WebUI.

The following procedure describes how to configure Remote AP static inner IP address:

To view IP address parameter in the local database:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Authentication** > **Auth Servers** tab.
2. In the **All Servers** table, select **Internal** under the **Name** field.
3. A list of **StaticIP for RAPs** is displayed under **Server > Internal** table.

To view the IP Address parameter in the Remote AP Whitelist:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Access Points > Remote APs** tab.

The following CLI commands configure Remote AP static inner IP address:

```
(host) [mynode] (config) #local-userdb add {generate-username|username <name>} {generate-password|password  <password>} {remote-ip <remote-ip>}
(host) [mynode] (config) #local-userdb modify {username < name>} {remote-ip <remote-ip>}
```

Issue the following command in config mode:

```
(host) [mynode] (config) #whitelist-db rap add {mac-address <address>}{ap-group <ap_group>}{remote-ip <remote-ip>}
(host) [mynode] (config) #whitelist-db rap modify {mac-address <address>} {remote-ip<remote-ip>}
```

Provisioning the AP

You need to configure the VPN client settings on the AP to instruct the AP to use IPsec to connect to the Managed Device. You can provision the Remote AP and allow remote users to provision the AP at home. This method of provisioning is referred as ZTP. See Provisioning 4G USB Modems on Remote Access Points on page 887 for more information about ZTP of remote AP.

You must provision the AP before you install it at its remote location. To provision the AP, the AP must be physically connected to the local network or directly connected to the Managed Device. 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 Managed Device.

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. In the **Managed Network** node hierarchy, navigate to the **Configuration > Access Points > Remote APs** tab.
2. Select the remote AP and click **Provision**.
3. Select **Pre-shared Key** from the **Authentication Methods** drop-down list.
4. Enter the **User name**, **Password**, and **Confirm password**.



The username and password you enter must match the username and password configured on the authentication server for the remote AP.

5. Select the **Static** option in the **Controller Discovery** field.
6. Set the **Controller IP/DNS** name as shown below:

Table 154: *Configuring a Managed device IP Address*

Deployment Scenario	Master IP Address Value
Deployment 1	Managed device IP address.
Deployment 2	Managed device public IP address.
Deployment 3	Public address of the NAT device to which the managed device is connected.



The username and password you enter must match the username and password configured on the authentication server for the remote AP.

7. Select **DHCP** option in the **IP** field.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Secondary Managed Device

The secondary managed device provides reliability and redundancy; however the functionality of a secondary managed device is initiated only after an AP terminates on a managed device successfully and retrieves the configuration. If the AP boots up and fails to connect to the managed device the AP cannot be managed. To address this, ArubaOS 8.0 introduces the secondary managed device feature.

In a scenario where the managed device is not reachable, the AP will try to reach the secondary managed device and if successful will terminate on the secondary managed device. The secondary managed device details are not stored in the system flash when the AP is deployed for the first time, but only after a successful configuration. An AP can use the secondary managed device feature after the AP reboots.



If an AP has not been configured to a managed device after deployment, the secondary managed device feature will not be applicable.

The following procedure describes how to enable the secondary managed device feature:

1. In the **Managed Network** node hierarchy, navigate to **Configuration** > **System** > **Profiles**.
2. Select **AP > AP System** under **All Profiles.**
3. Select the AP profile for which the secondary managed device feature is to be enabled. The AP system profile section is displayed.
4. Enter an IP or FQDN value for the secondary managed device in the **Secondary Master IP/FQDN** field.
5. Click **Save**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.



The secondary managed device feature can be enabled on the secondary managed device.

The following CLI commands enable the secondary managed device feature.

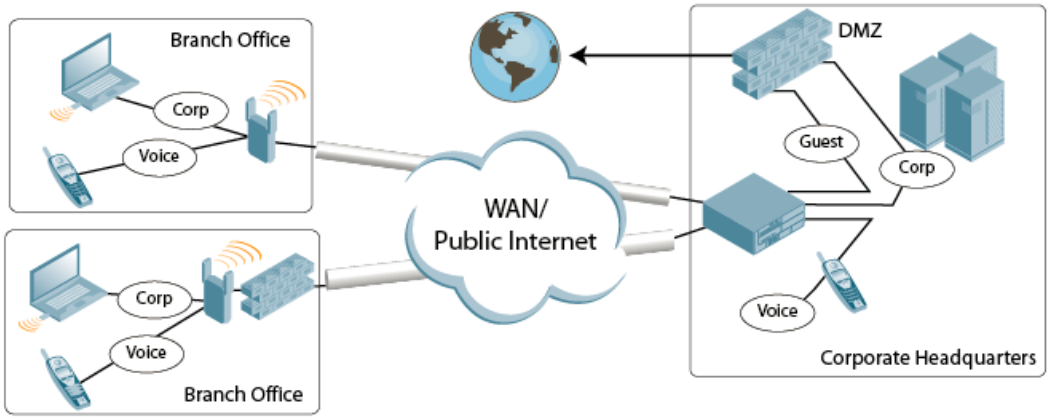
```
(host) [mynode] (config) #ap system-profile <profile name>
(host) [mynode] (AP system profile "profile name")#secondary-master <value>
```

Deploying a Branch or Home Office Solution

In a branch office, the AP is deployed in a separate IP network from the corporate network. Typically, there are one or two NAT devices between the two networks. Branch office users need access to corporate resources such as printers and servers, but traffic to and from these resources must not impact the corporate head office.

Figure 102 is a graphic representation of a remote AP in a branch or home office, with a single Managed Device providing access to both a corporate WLAN and a branch office WLAN.

Figure 102 *Remote AP with Single Managed device*



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 Managed Device is used as a RADIUS pass-through when the authenticator has to communicate with a RADIUS server (which also supports survivability).
- 802.11 encryption or decryption is in the AP to provide access to local resources.
- Local bridging of client traffic connected to the WLAN or to an AP enet1 port to provide access to local resources.

Provisioning the Branch AP

You can provision the remote AP either using the Managed Device or using the ZTP method. For more information on Managed Device provisioning, see Configuring Installed APs on page 643. For more information on ZTP, see Provisioning 4G USB Modems on Remote Access Points on page 887.

Configuring the Branch AP

- Specify forward mode for the ESSID in the virtual AP profile
- Specify remote AP operation in the virtual AP profile (The remote AP operates in standard mode by default.)
- Set how long the AP stays up after connectivity to Managed Device has gone down in the SSID profile
- Set the VLAN ID in the virtual AP profile
- Set the native VLAN ID in the AP system profile
- Set forward mode for enet1 port



Remote APs support 802.1q VLAN tagging. Data from the remote AP will be tagged on the wired side.

Troubleshooting Remote AP

The following WebUI options are available to troubleshoot issues with remote AP:

- Using local debugging feature
- Viewing the remote AP summary report
- Viewing remote AP connectivity report
- Using remote AP diagnostic options

Local Debugging

Local debugging is a WebUI feature that allows end users to perform diagnostics and view the status of their remote AP through a wired or wireless client. This feature is useful for troubleshooting connectivity problems on remote APs and performing throughput tests. There are three tabs in the **Local Debugging** WebUI window; **Summary**, **Connectivity**, and **Diagnostics**. Each tab displays different information for the AP, but all three tabs include a **Generate & save support file** link that, when clicked, will automatically generate a **support.tgz** file that can be sent to a corporate IT department for additional analysis and debugging.



A snapshot of the bridge, acl, session, user, and arp tables, current processes, memory, and kernel debug messages are captured in a single **rap_debug.txt** file which is bundled along with **support.tgz** file.

Remote AP Summary

The **Summary** tab has two views; basic and advanced. Click the **basic** or **advanced** links at the top of this tab to toggle between the two views. The table below shows the information displayed for both the basic and advanced views of the **Summary** tab.

Table 155: *Remote AP Console Summary Tab Information*

Summary Table Name	Basic View Information	Advanced View Information
Wired Ports Status	▪ **Port**: port numbers of the wired ports on the AP ▪ **Status**: current status of each port (*Connected*, *LinkDown* or *Disabled*).	The advanced view of the Wired Access Ports table displays the following data: ▪ **Port**: port numbers of the wired ports on the AP ▪ **Status**: current status of each port (*Connected*, *LinkDown* or *Disabled*) ▪ **MAC Address:** MAC address of the wired port ▪ **Speed**: speed of the link ▪ **Duplex Type**: duplex mode of the link, full or half ▪ **Forwarding mode**: forwarding mode for the port: *Bridge*, *Decrypt-Tunnel*, *Split-Tunnel*, or *Tunnel* ▪ **Users**: fumber of users accessing each port ▪ **Rx Packets**: number of packets received on the port ▪ **Tx packets**: number of packets transmitted via the port
Wireless SSIDs	▪ **SSID**: Name of the SSID. ▪ **Status**: SSID Status (up, down, or disabled). ▪ **Band**: Radio band available on the SSID.	▪ **SSID**: name of the SSID ▪ **Status**: SSID Status (up, down, or disabled). ▪ **Band**: radio band available on the SSID ▪ **Channel**: channel used on the radio band ▪ **BSSID**: BSSID of the wireless SSID ▪ **Forwarding Mode**: forwarding mode used by the Wireless SSID (*Bridge*, *Decrypt-Tunnel*, *Split-Tunnel*, or *Tunnel*) ▪ **EIRP**: equivalent Isotropic Radiated Power, in dBm ▪ **Noise floor**: residual background noise detected by an AP. Noise seen by an AP is reported as -dBm Therefore, a noise floor of -100 dBm is smaller (lower) than a noise floor of -50 dBm. ▪ **Users**: number of users on the radio

Summary Table Name	Basic View Information	Advanced View Information
		band ▪ **Rx Packets**: number of packets received on the BSSID ▪ **Tx packets**: number of packets transmitted via the BSSID
Wired Users	▪ **MAC Address**: MAC address of the wired user. ▪ **IP address**: IP address of the wired user.	▪ **MAC Address**: MAC address of the wired user. ▪ **IP address**: IP address of the wired user. ▪ **Port**: AP port used by the wired user.
Wireless User	▪ **MAC Address**: MAC address of the wireless user. ▪ **IP address**: IP address of the wireless user.	▪ **MAC Address**: MAC address of the wired user ▪ **IP address**: IP address of the wired user ▪ **SSID**: name of the SSID ▪ **BSSID**: BSSID of the wireless user ▪ **Assoc State**: shows if the user is associated or just authorized ▪ **Auth**: Type of authentication: WPA, 802.1X, none, open, or shared ▪ **Encryption**: encryption type used by the wireless user ▪ **Band**: radio band used by the wireless client ▪ **RSSI**: Receive Signal Strength Indicator (RSSI) value displayed in the output of this command represents signal strength as a signal to noise ratio.
Device Info	▪ **Type**: AP device or 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 managed device. ▪ **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 Mobility	N/A

Summary Table Name	Basic View Information	Advanced View Information
	Master. ■ **Ims**: IP address of the managed device.	
Uplink Info	The Uplink Info table can display some or all of the following information for your remote AP, depending upon whether a link is active and the number of links supported by the AP. Active uplink information, including: ■ Interface name ■ Port speed ■ IP address Standby link information, including: ■ Name (3G) ■ **Device connected** (yes or no) ■ **Provisioned** (yes or no) ■ IP address ■ Device ■ User ■ Password	N/A

Multihoming on Remote AP

You can uplink a Remote AP as an Ethernet or a USB based modem. These uplinks can be used as a backup link if the primary link fails. The uplink becomes active based on the order of priority configured on the Remote AP. The Remote AP switches back to the primary link when the primary connection is restored.

For information on provisioning the Remote AP using the USB based modem, see Provisioning 4G USB Modems on Remote Access Points on page 887.

Seamless failover from backup link to primary link on Remote AP

Remote APs can failover from a backup link to a primary link without much disruption to traffic. Also the failover is performed only if the Managed Device is reachable via the primary link.

Remote AP Connectivity

The information shown on the **Connectivity** tab will vary, depending upon the current status of the Remote AP. If a remote AP has been successfully provisioned and connected, it should display some or all of the information in Table 156.

Table 156: *Remote AP 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.

Data	Description
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 TPM certificate.
Master Connectivity	Shows if the AP was able to connect to the Managed Device. This item also shows the IP address to which the AP attempted to connect, and, if the AP did connect successfully, the link used to connect to that Managed Device.
LMS Connectivity	Shows if the AP was able to connect to a managed device. This item also shows the IP address to which the AP attempted to connect, and, if the AP did connect successfully, the link used to connect to that Managed Device.

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.

To run a diagnostic test on a remote AP:

1. In the **Managed Network** node hierarchy, navigate to the **Diagnostics> Tools** menu.
2. Select **Ping**, **Traceroute**, or**Tracepath** tab.

   The *ping* and *traceroute* tests require that you enter a network destination in the form of an IP address or FQDN, and select either **bridge** or **tunnel** mode for the test. *The NSLookup* diagnostic test requires that you enter a destination only. The test checks the link between the AP and the Mobility Master, and does not require any additional test configuration settings.
3. Click **Trace** to start the test. The results of the test will appear in the **Diagnostics** window.
4. To view log files in a separate browser window:
   a. Click **Diagnostics> Logs** menu.
   b. Select a log file name from the **Logs** drop-down list. The type of log files available will vary, depending upon your remote AP configuration.
   c. Click **Display**.

Bringing up Certificate-Based Remote AP in VMC

A certificate-based remote AP does not come up on a virtual mobility controller (VMC) because TPM certificate for the AP is present in the Mobility Master. However, you can bring up the remote AP by using a self-signed certificate.

To do this, first you need to bring up the AP as campus AP. Then, reprovision the AP to come up as remote AP.

The following procedure describes how to bring up a remote AP on the VMC by using a self-signed certificate:

1. Bring up the AP as a campus AP. See the Configuring Installed APs section.



Before reprovisioning the AP as remote AP, ensure that the AP has come up as campus AP successfully.

2. In the **Managed Network** node hierarchy, navigate to the **Configuration > Access Points > Campus APs** tab.
3. Select the AP name you want to reprovision as remote AP.
4. Click **Provision**.
5. Select **Static** option in the **Controller discovery** field.
6. Enter the IP address or the complete DNS name for the Managed Device, In the **Controller IP/DNS name** field,
7. Select **DHCP** option in the **IP** field.
8. Select the **Remote** option in the **Deployment** field.
9. Select **Certificate** from the **Authentication methods** drop-down list.
10. Select **self-signed** from the **Trust anchor** drop-down list.
11. Click **Submit**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Hybrid Model Support for Remote AP Terminating on a VMC

The hybrid model of deploying Remote APs addresses the following issues:

- Self-signed certificates used by the VMC has a 10-year validity, which does not cater to users who do not prefer using self-signed certs with more than 1-year validity.
- Users are unable to move Remote APs from one VMC to another as the two VMCs use different self-signed certs, due to which RAPs fail to establish a tunnel with new VMC.

In the hybrid model these issues are mitigated by deploying the following implementation:

- Remote APs will use factory certificates.
- VMCs will use custom certificates.

Prerequisites

- The user should load custom CA and server certs on the VMC.
- Custom CA and server certs certificate group should be created.
- Push custom CA to the Remote AP.

Moving a Remote AP

Use the following steps to move a Remote AP from one VMC to another in the Hybrid model:

1. Load the custom CA of VMC2 on VMC1.
2. Move the Remote AP from VMC1 to VMC2 by pushing the custom CA of VMC2.



If you are using a custom server certificate signed by same custom CA on VMC1 and VMC2, then RAP moving between VMC's doesn't need any further certificate push from any VMC.

Enabling Remote AP Advanced Configuration Options

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

- Understanding Remote AP Modes of Operation on page 859
- Working in Fallback Mode on page 861
- Specifying the DNS Managed Device Setting on page 869
- Backup Managed Device List on page 870
- Configuring Remote AP Failback on page 871
- Working with ACL and Firewall Policies on page 873
- Understanding Split Tunneling on page 873
- Provisioning Wi-Fi Multimedia on page 886



The information in this section assumes you have already configured the remote AP functionality, as described in Configuring the Secure Remote Access Point Service on page 845.

Understanding Remote AP Modes of Operation

Table 157 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 Managed Device using GRE, bridged to the local Ethernet LAN, or a combination thereof) and the remote AP mode of operation (when the virtual AP operates on a remote AP) in the virtual AP profile.

The column on the left of the table lists the remote AP operation settings. The row across the top of the table lists the forward mode settings. To understand how these settings work in concert, scan the desired remote AP operation with the forward mode setting, and read the information in the appropriate table cell.

The all column and row lists features that all remote AP operation and forward mode settings have in common regardless of other settings. For example, at the intersection of all and bridge, the description outlines what happens in bridge mode regardless of the remote AP mode of operation.

Table 157: *Remote AP Modes of Operation and Behavior*

Remote AP Operation Setting	Forward Mode Setting				
	all	**bridge**	**split-tunnel**	**tunnel**	**decrypt-tunnel**
all		Management frames on the AP. Frames are bridged between wired and wireless interfaces. No frames are tunneled to the Managed Device.	Management frames on the AP.	Frames are GRE tunneled to the Managed Device to an untrusted tunnel. 100% of station frames are tunneled to the Managed Device.	Management frames on the AP. Frames are always GRE tunneled to Managed Device.

Remote AP Operation Setting	Forward Mode Setting				
		Station acquires its IP address locally from an external DHCP server.	Frames are either GRE tunneled to the Managed Device, to a trusted tunnel or are sent through the NAT and bridged on the wired interface according to user role and session ACL. Typically, the station obtains an IP address from a VLAN on the Mobility Master. Typically, the AP has ACLs that forward corporate traffic through the tunnel and source NAT the non-corporate traffic to the Internet.		
always	ESSID is always up when the AP is up regardless of whether the Managed Device is reachable. Supports PSK ESSID only. SSID configuration stored in flash on AP.	Provides an SSID that is always available for local access.	Not supported	Not supported	Not supported
	all	**bridge**	**split-tunnel**	**tunnel**	**decrypt-tunnel**
backup	ESSID is only up when the Managed Device is unreachable. Supports PSK ESSID only. SSID configuration stored in flash on AP.	Provides a backup SSID for local access only when the Managed Device is unreachable.	Not supported	Not supported	Not supported

Remote AP Operation Setting	Forward Mode Setting				
persistent	ESSID is up when the AP contacts the Managed Device and stays up if connectivity is disrupted with the Managed Device. SSID configuration obtained from the Managed Device. Designed for 802.1X SSIDs.	Same behavior as standard, described below, except the ESSID is up if connectivity to the Managed Device is lost.	Not supported	Not supported	Not supported
standard	ESSID is up only when there is connectivity with the Managed Device. SSID configuration obtained from the Managed Device.	Behaves like a classic Aruba branch office AP. Provides a bridged ESSID that is configured from the Managed Device and stays up if there is Managed Device connectivity.	Split tunneling mode	Classic Aruba thin AP operation	Decrypt tunnel mode

Working in Fallback Mode

The fallback mode (also known as backup configuration) operates the remote AP if the master Managed Device or the configured primary and backup LMS are unreachable. The remote AP saves configuration information that allows it to operate autonomously using one or more SSIDs in local bridging mode, while supporting open association or encryption with PSKs. You can also use the backup configuration if you experience network connectivity issues, such as the WAN link or the central data center becoming unavailable. With the backup configuration, the remote site does not go down if the WAN link fails or the data center is unavailable.

You define the backup configuration in the virtual AP profile on the Managed Device. The remote AP checks for configuration updates each time it establishes a connection with the Managed Device. 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 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 Managed Device. This SSID is advertised until the Managed Device is reachable. Recommended for bridge SSIDs.
- Persistent - Permanently enables the virtual AP after the remote AP initially connects to the Managed Device. Recommended for 802.1X SSIDs.
- Standard - Enables the virtual AP when the remote AP connects to the Managed Device. 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 Managed Device. If you configure the remote AP in backup mode, and a connection to the Managed Device 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 Managed Device has been re-established.

Backup Configuration Behavior for Wired Ports

If the connection between the remote AP and the Managed Device is disconnected, the remote AP will be exhibit the following behavior:

- All access ports on the remote AP will be moved to bridge forwarding mode ,irrespective of their original forwarding mode.
- Clients will receive an IP address from the remote AP's DHCP server.
- Clients will have complete access to Remote AP's uplink network. You cannot enforce or modify any access control policies on the clients connected in this mode.

This section describes the following topics:

- Configuring Fallback Mode on page 862
- Configuring the DHCP Server on the Remote AP on page 863
- Configuring Advanced Backup Options on page 865

Configuring Fallback Mode

To configure the fallback mode, you must:

- Configure the AAA profile
- Configure the virtual AP profile

Configuring the AAA Profile for Fallback Mode

The following procedure describes how to configure the AAA profile for fallback mode:

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

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Select **Wireless LAN > AAA** under **All Profiles.**
3. In the **AAA Profile: New Profile** table, click **+** in **AAA Profile**.
4. Enter the **Profile name**.
5. For **Initial role**, select the appropriate role (for example, logon) from the drop-down list.
6. For **802.1X Authentication Default Role**, select the appropriate role (for example, default) from the drop-down list.
7. Click **Save**.
8. Select the AAA profile that you just created:
    a. Click **802.1X Authentication Server Group**, and select the **Server Group** to be used (for example, default) from the drop-down list.

b. Click **Save**.



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.

c. Click **802.1X Authentication**, and select the **802.1X Authentication Profile** to be used (for example, default) from the drop-down list.

d. Click **Save**.

9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the AAA profile for fallback mode:

```
(host) [md] (config) #aaa profile default
(host) [md] (AAA Profile "default") #initial-role <role>
(host) [md] (AAA Profile "default") #authentication-dot1x <dot1x-profile>
(host) [md] (AAA Profile "default") #dot1x-default-role <role>
(host) [md] (AAA Profile "default") #dot1x-server-group <group>
```

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 Managed Device 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:

1. 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.
2. Specify the DHCP IP address pool and netmask. The AP assigns IP addresses from the DHCP pool 192.168.11.0/24 by default, with an IP address range from 192.168.11.2 through 192.168.11.254. You can manually define the DHCP IP address pool and netmask based on your network design and IP address scheme.
3. Specify the IP address of the DHCP server, DHCP router, and the DHCP DNS server. The AP uses IP address 192.168.11.1 for the DHCP server, the DHCP router, and the DHCP DNS server by default.
4. Enter the amount of days the assigned IP address is valid (also known as the remote AP DHCP lease). The lease does not expire by default, which means the IP address is always valid.
5. Assign the VLAN ID for the remote AP DHCP VLAN to a virtual AP profile. When a client connects to that virtual AP profile, the AP assigns the IP address from the DHCP pool.



The following is a high-level description of the steps required to configure the DHCP server on the remote AP. The steps assume you have already created the virtual AP profile, AAA profile, SSID profile, and other settings for your remote AP operation (for information about the backup configuration, see Configuring the DHCP Server on the Remote AP on page 863).

The following procedure describes how to configure the DHCP Server on the Remote AP:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > AP Groups** menu option.
2. Select an **AP Group**.
3. Click **More** tab and add the appropriate value in the following fields:
   a. Enter the LMS IP address in the **LMS IP address** field and the Bakcup LMS IP address in the **Backup LMS IP address** field.
   b. Enter the LMS IPv6 address in the **LMS IPv6 address** field and the Bakcup LMS IPv6 address in the **Backup LMS IPv6 address** field.
   c. Click **Submit**.
4. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles** tab.
   a. Select **AP> AP System** under **All Profiles**, and select an AP system profile.
   b. Click **Remote AP** accordion.
   c. Enter the VLAN ID of the backup configuration virtual AP VLAN in the **Remote-AP DHCP Server VLAN** field .
   d. Enter the IP address of the DHCP server in the **Remote-AP DHCP Server ID** field.
   e. Enter the IP address of the default DHCP router in the **Remote-AP DHCP Default Router** field.
   f. Specify the DHCP IP address pool. This configures the pool of IP addresses from which the remote AP uses to assign IP addresses.
      i. Enter the first IP address of the pool, in the **Remote-AP DHCP Pool Start** field.
      ii. Enter the last IP address of the pool, in the **Remote-AP-DHCP Pool End** field.
      iii. Enter the netmask, in the **Remote-AP-DHCP Pool Netmask** field.
   g. Specify the number of days for which the IP address is valid, in the **Remote-AP DHCP Lease Time** field.
   h. Click **Save**.
   i. Select **Wireless LAN > Virtual AP** under **All Profiles**, and the select virtual AP profile you want to configure.
   j. Click **Save**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the DHCP Server on the Remote AP:

```
(host) [md] (config) #ap system-profile default
(host) [md] (AP system profile "default") #lms-ip <ipaddr>
(host) [md] (AP system profile "default") #master-ip <ipaddr>
(host) [md] (AP system profile "default") #rap-dhcp-default-router <ipaddr>
(host) [md] (AP system profile "default") #rap-dhcp-dns-server <ipaddr>
(host) [md] (AP system profile "default") #rap-dhcp-lease <days>
(host) [md] (AP system profile "default") #rap-dhcp-pool-end <ipaddr>
(host) [md] (AP system profile "default") #rap-dhacp-pool-netmask <netmask>
(host) [md] (AP system profile "default") #rap-dhcp-pool-start <ipaddr>
(host) [md] (AP system profile "default") #rap-dhcp-server-id <ipaddr>
(host) [md] (AP system profile "default") #rap-dhcp-server-vlan <vlan>
```

```
(host) [md] (config) #wlan virtual-ap default
(host) [md] (Virtual AP profile "default") #ssid-profile <profile>
(host) [md] (Virtual AP profile "default") #vlan <vlan>
(host) [md] (Virtual AP profile "default") #forward-mode bridge
(host) [md] (Virtual AP profile "default") #aaa-profile <name>
(host) [md] (Virtual AP profile "default") #rap-operation {always|backup|persistent}

(host) [md] (config) #ap-group default
(host) [md] (AP group "default") #virtual-ap <name>
```

or

```
(host) [md] (config) #ap-name default
(host) [md] (AP name "default") #virtual-ap <name>
```

Configuring Advanced Backup Options

You can also use the backup configuration (fallback mode) to allow the remote AP to pass through a captive portal, such as network access in a hotel, airport, or other public network, to access the corporate network. For this scenario:

- Define a session ACL for the bridge SSID to source NAT all user traffic, except DHCP. For example, use **any any svc-dhcp permit** followed by **any any any route src-nat**. Apply the session ACL to a remote AP user role.
- Configure the AAA profile. Make sure the initial role contains the session ACL previously configured.

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

  802.1X and PSK authentication is supported when configuring bridge or split tunnel modes.

- Configure the virtual AP profile for the backup configuration:
  - Set the remote AP operation to **always** or **backup**.
  - Create and apply the applicable SSID profile.
  - Configure a bridge SSID for the backup configuration. In the virtual AP profile, specify forward mode as **bridge**.

    For more information about the backup configuration, see Configuring Fallback Mode on page 862.
- 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 863.

  If you use a local DHCP server to obtain IP addresses, you must define one additional ACL to permit traffic between clients without using source NAT to route the traffic. Using the previously configured ACL, add **user alias internal-network any permit** before **any any any route src-nat**.
- Connect the remote AP to the available public network (for example, a hotel or airport network).

  The remote AP advertises the backup SSID so the wireless client can connect and obtain an IP address from the available DHCP server.

  The client can obtain an IP address from the public network, for example a hotel or airport, or from the DHCP server on the remote AP.

  After obtaining an IP address, the wireless client can connect and access the corporate network and bring up the configured corporate SSIDs.

The following is a high-level description of what is needed to configure the remote AP to pass through a captive portal and access the corporate Managed Device. This information assumes you are familiar with configuring session ACLs, AAA profiles, virtual APs, and AP system profiles and highlights the modified parameters.

Configuring the Session ACL

The following procedure describes how to configure session ACL:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration** >**Roles and Policies** >**Policies** tab.
2. Click **+** to create a new policy.
3. Enter the policy name in the **Policy Name** field.
4. From the **Policy Type** drop-down list, select **Session**.
5. Click **Submit**.
6. To create the first rule:
    a. Select the policy created.
    b. Click **+** in the **Policies > New Policy** table.
    c. Select **Access Control** option in the **Rule Type** field.
    d. Click **OK**.
    e. Select **Any** from the **Source** drop-down list.
    f. Select **Any** from the **Destination** drop-down list.
    g. Select **Service** from the **Service/app** drop-down list.
    h. Select **svc-dhcp** from the **Service alias** drop-down list.
    i. Select **Permit** from the **Action** drop-down list.
    j. Click **Submit**.
7. To create the next rule:
    a. Click the policy created.
    b. Click **+**.
    c. Select **Access Control** option in the **Rule Type** field.
    d. Click **OK**.
    e. Select **Any** from the **Source** drop-down list.
    f. Select **Any** from the **Destination** drop-down list.
    g. Select **Service** from the **Service/app** drop-down list.
    h. Select **Any** from the **Service alias** drop-down list.
    i. Select **Route Source NAT** from the **Action** drop-down list.

8. Click **Submit**.



If you use a local DHCP server to obtain IP addresses, you must define one additional ACL to permit traffic between clients without using source NAT to route the traffic. Add `user alias internal-network any permit` before `any any any route src-nat`.

9. In the **Managed Networks** node hierarchy, navigate to the **Configuration** >**Roles and Policies** >**Roles** tab.



Roles can be created only in the managed device.

10. Click **+** to create a new role.
11. Enter the role name in the **Name** field.
12. Click **Submit**.
13. Select the new role created.
14. Click **Show Advanced View**.
15. Click **+**.
16. Select an **Add existing policy** option and select the policy created from the **Policy name** drop-down list.
17. Click **Submit**.
18. Click **Pending Changes**.
19. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure session ACL:

```
(host) [md] (config) #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 using source NAT to route the traffic. Add **user alias internal-network any permit** before **any any any route src-nat:**

```
(host) [md] (config) #user-role <role>
  session-acl <policy>
```

Configuring the AAA Profile

The following procedure describes how to configure the AAA profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Select **Wireless LAN > AAA** under **All Profiles.**
3. Click **+** in **AAA Profile**.
   a. Enter the **Profile name**, and select the AAA profile that you just created:
   b. For **Initial role**, select the user role you just created.
   c. For **802.1X Authentication Default Role**, select the appropriate role (for example, “default”) from the drop-down list.
   d. Click **Save**.

e. Under the AAA profile that you created, locate **802.1X Authentication Server Group**, and select the Server Group to be used (for example default) from the drop-down list.

4. Click **Save**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the AAA profile:

```
(host) [md] (config) #aaa profile default
(host) [md] (AAA Profile "default") #initial-role <role>
```

Defining the Backup Configuration

The following procedure describes how to define the backup configuration:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Select **Wireless LAN > Virtual AP** under **All Profiles.**
3. Click **+** in **Virtual AP profile** and enter the profile name.
4. Click **Save**.



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. Click **+** next to the profile created.

b. Click **AAA** and select a previously configured AAA profile from the **AAA Profile** drop-down list.

c. Click **Save.**

d. Click **SSID** and select the previously configured SSID profile from the **SSID Profile** drop-down list.

e. Click **Save.**

5. Select the new virtual AP name listed under **Wireless LAN > Virtual AP**, to view the configuration parameters.
6. In the **General** accordion under **Virtual AP profile: name**, execute the following:
   a. Ensure that the **Virtual AP enable** check box is selected.
   b. Enter the VLAN ID to be used for the Virtual AP profile in the **VLAN** field.
   c. Select **bridge** from the **Forward mode** drop-down list.
   d. Click **Save**.
7. Under **All Profiles**, select **AP>AP system** profile.
8. Under **Profile Details**, execute the following:
   a. Select the AP system profile that you want to edit.
   b. Under the **LMS Settings** accordion, enter the LMS IP address in the **LMS IP** field.
   c. Under the **Remote AP** accordion, enter the Remote -AP DHCP server name in the **Remote-**

AP DHCP Server field.

d. Click **Save**.

9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands define the backup configuration:

```
(host) [mynode] (config) #wlan ssid-profile <profile>
essid <name>
opmode <method>
wpa-passphrase <string> (if necessary)

(host) [mynode] (config) #wlan virtual-ap <name>
ssid-profile <profile>
vlan <vlan>
forward-mode bridge
aaa-profile <name>
rap-operation {always|backup}

(host) [mynode] (config) #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>

(host) [mynode] (config) #ap-group <name>
 virtual-ap <name>
 ap-system-profile <name>
```

or

```
(host) [mynode] (config) #ap-name <name>

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

Specifying the DNS Managed Device Setting

In addition to specifying IP addresses for Managed Device, you can also specify the master DNS name for the Managed Device when provisioning the remote AP. The name must be resolved to an IP address when attempting to set up the IPsec tunnel. For information on how to configure a host name entry on the DNS server, refer to the vendor documentation for your server. It is recommended to use a maximum of 8 IP addresses to resolve a Managed Device 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 Managed Device. For more detailed information, see the next section Backup Managed Device List .

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 Managed Device 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 Managed Device information in the Master Discovery section of the Provision page.



Reprovisioning the AP causes it to automatically reboot.

The following procedure describes how to specify the DNS managed device setting:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Access Points > Remote APs** tab.
2. Select the remote AP and click **Provision**.
3. In the **General** tab, enter master DNS name in the **Controller IP/DNS name** field.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

For more information, see Configuring the Secure Remote Access Point Service on page 845.

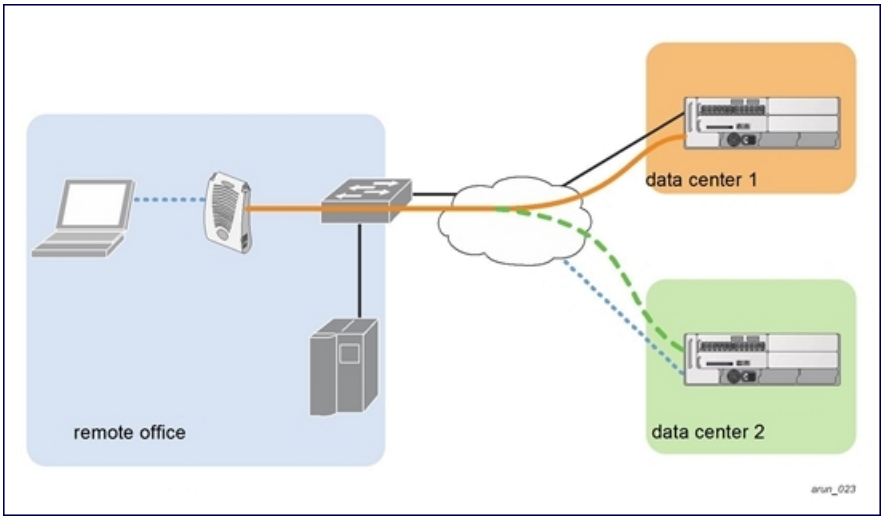
Backup Managed Device List

Using DNS, the remote AP receives multiple IP addresses in response to a host name lookup. Known as the backup managed device list, remote APs go through this list to associate with a managed device. If the primary managed device is unavailable or does not respond, the remote AP continues through the list until it finds an available managed device. This provides redundancy and failover protection.

The remote AP loses the IP address information received through DNS when it terminates and receives the system profile configuration from the managed device. If the remote AP loses connectivity on the IPsec tunnel to the managed device, the Remote AP fails over from the primary managed device to the backup managed device. For this scenario, add the IP address of the backup managed device in the backup LMS and the IP address of the primary managed device in the LMS field of the ap-system profile. Network connectivity is lost during this time. As described in the section Backup Managed Device List on page 870, you can also configure a remote AP to revert back to the primary managed device 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 managed device in the DMZ. You can provision the remote APs to use the managed device in data center 1 as the primary managed device, and the managed device in data center 2 as the backup managed device. 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 103 Sample Backup Scenario



Configuring the LMS and Backup LMS IP Addresses

The following procedure describes how to configure the LMS and Backup LMS IP addresses:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Select **AP > AP system** under **All Profiles.**
3. Select the AP system profile you want to modify.
4. Enter the following details in the **LMS Settings** accordion:
   a. Enter the primary managed device IP address, in the **LMS IP** field.
   b. Enter the backup managed device IP address, in the **Backup LMS IP** field.
5. Click **Save**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configures the LMS and Backup LMS IP addresses:

```
(host) [md] (config) #ap system-profile <profile>
lms-ip <ipaddr>
bkup-lms-ip <ipaddr>

(host) [md] (config) #ap-group <group>
ap-system-profile <profile>

(host) [md] (config) #ap-name <name>
ap-system-profile <profile>
```

Configuring Remote AP Failback

In conjunction with the backup managed device list, you can configure remote APs to revert back (failback) to the primary managed device if it becomes available. If you do not explicitly configure this behavior, the remote AP will keep its connection with the backup managed device until the remote AP, managed device, 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 managed device list and attempt to connect with the primary managed device.

The following procedure describes how to configure Remote AP failback:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Select **AP > AP system** under **All Profiles.**
3. Select the AP system profile you want to modify.
4. Enter the following details in the **LMS Settings** accordion:
   a. Select the **LMS Preemption** check box. This is disabled by default.
   b. Enter the duration (in seconds) for which the remote AP must wait before moving back to the primary managed device, in the **LMS Hold-down period** field.
5. Click **Save**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures Remote AP failback:

```
(host) [md] (config) #ap system-profile <profile>
lms-preemption
lms-hold-down period <seconds>
```

Enabling Remote AP Local Network Access

You can enable local network access between the clients (from same or different subnets and VLANs) connected to a Remote AP through wired or wireless interfaces in split-tunnel or bridge forwarding modes. This allows the clients to effectively communicate with each other without routing the traffic via the managed device. You can use WebUI or CLI to enable the local network access.

The following procedure describes how to enable Remote AP local network access:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Select **AP > AP system** under **All Profiles.**
3. Select the AP system profile you want to modify.
4. To enable remote network access, select the **Remote-AP Local Network Access** check box under the **Remote AP** accordion.
5. Click **Save**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

- The following CLI commands enable Remote AP local network access:

```
(host) [md] (config) #ap system-profile <ap-profile> rap-local-network-access
```

- To disable, enter the following command:

```
(host) [md] (config) #ap system-profile <ap-profile> no rap-local-network-access
```

See the *ArubaOS CLI Reference Guide* for detailed information on the command options.

Configuring Remote AP Authorization Profiles

Remote AP configurations include an authorization profile that specifies which profile settings should be assigned to a remote AP that has been provisioned but not yet authenticated at the remote site. These yet-unauthorized APs are put into the temporary AP group **authorization-group** by default and

assigned the predefined profile **NoAuthApGroup**. This configuration allows the user to connect to an unauthorized remote AP via a wired port, then enter a corporate username and password. Once a valid user has authorized the AP, and it will be marked as authorized on the network. The remote AP will then download the configuration assigned to that AP by its permanent AP group.

The following procedure describes how to create a new authorization profile or edit an existing authorization profile:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Select **AP > AP Authorization** under **All Profiles.**
   a. To edit an existing profile, select a profile listed under **AP Authorization** profile and select a new AP authorization group from the **AP authorization group** drop-down list.
   b. To create a new authorization profile, click **+** next to the **AP Authorization profile** field.
      - Enter the name in the **Profile name** field.
      - Select a group from the **AP authorization group** drop-down list.
3. Click **Save**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands create a new authorization profile or edit an existing authorization profile:

```
(host) [md] (config) #ap authorization-profile <profile>
authorization-group <ap-group>
```

Working with ACL and Firewall Policies

Remote APs support the following ACL; 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 Aruba Managed Device and takes some action based on that identification. You apply these ACLs to user roles or uplink ports.



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

For more information about ACLs and firewall policies, see Working with ACL and Firewall Policies on page 873.

Understanding Split Tunneling

The split tunneling feature allows you to optimize traffic flow by directing only corporate traffic back to the Managed Device, while local application traffic remains local. This ensures that local traffic does not incur the overhead of the round trip to the Managed Device, 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 Managed Device and local traffic.

Figure 104 Sample Split Tunnel Environment

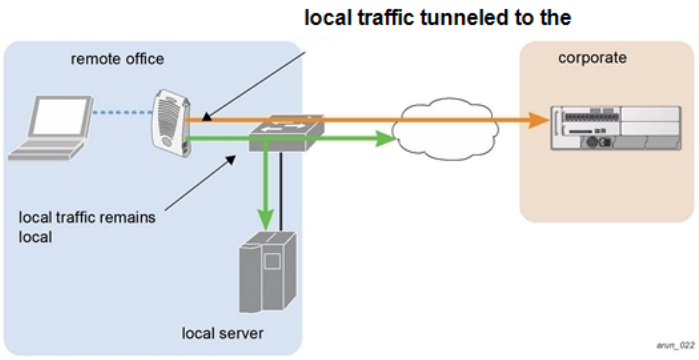


Figure 104 displays corporate traffic which is GRE tunneled to the Managed Device through a trusted tunnel and local traffic is sent through the source NAT and bridged on the wired interface based on the configured user role and session ACL.

Configuring Split Tunneling

The procedure to cconfigure split tunneling requires the following steps. Each step is described in detail later in this chapter.



The split tunneling feature requires the PEFNG license. If you do not have the PEFNG license on your Managed Device, you must install it before you configure split tunneling. For details on installing licenses, refer to the *Aruba Managed Device Licensing Guide.*

1. Define a session ACL that forwards only corporate traffic to the Managed Device.
2. Configure a net destination for the corporate subnets.
3. Create rules to permit DHCP and corporate traffic to the corporate Managed Device.
4. Apply the session ACL to a user role.
5. (Optional) Configure an ACL that restricts remote AP users from accessing the remote AP local debugging homepage.
6. Configure the remote AP's AAA profile.
7. Specify the authentication method (**802.1X** or **PSK**) and the default user role for authenticated users. The user role specified in the AAA profile must contain the session ACL defined in the previous step.
8. (Optional) Use the remote AP's AAA profile to enable RADIUS accounting.
9. Configure the virtual AP profile:
10. Specify which AP group or AP to which the virtual AP profile applies.
11. Set the VLAN used for split tunneling. Only one VLAN can be configured for split tunneling; VLAN pooling is not allowed.
12. When specifying the use of a split tunnel configuration, use "split-tunnel" forward mode.

13. Create and apply the applicable SSID profile.



When creating a new virtual AP profile in the WebUI, you can also configure the SSID at the same time. For information about AP profiles, see AP Configuration Profiles on page 639.

14. (Optional) Create a list of network names resolved by corporate DNS servers.

Configuring the Session ACL Allowing Tunneling

First you need to configure a session ACL that "permits" corporate traffic to be forwarded (tunneled) to the Mobility Master, and that routes, or locally bridges, local traffic.

The following procedure describes how to configure the session ACL allowing tunneling:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Roles and Policies** > **Policies** tab.
2. Click **+** to create a new policy.

   The **New Policy** window is displayed.
3. Select **Session** from the **Policy Type** drop-down list.
4. Enter the policy name in the **Policy Name** field.
5. Click **Submit**.
6. To create the first rule:
   a. Select the policy created in the previous steps.
   b. Click **+**in the **Policy > <policy name> Rules** table .

      The **New Rule for <policy name>** window is displayed.
   c. Select **Access control** or **Application** in the **Rule Type**field.
   d. Click **OK**.
7. In the **New Policy > New Forwarding Rule** table, configure the following parameters:
   - **IP version**—Select **IPv4** or **IPv6** from the drop-down list.
   - **Source**—Select **Any** from the drop-down list.
   - **Destination**—Select **Any** from the drop-down list.
   - **Service/app**—Select **Service** from the drop-down list.
   - **Service alias**—Select **svc-dhcp** from the drop-down list.
   - **Action**—Select **Permit** from the drop-down list.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

    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.
11. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Roles and Policies** > **Aliases** tab.
12. Click **+** In the **Network Aliases** pane.
13. In the **Destination** pane, configure the following parameters:
    - **IP version**—Select **IPv4** or **IPv6** from the drop-down list.
    - **Destination name**—Enter a name.

- **Description**—Enter a description of the destination within 128 characters.
- **Invert**—Select the check box to specify that the inverse of the network addresses configured are used.

14. Under **Items**, click **+**.

    The **Add New Destination** window is displayed.

15. Configure the following parameters:
    - **Rule Type**—Select **Network** from the drop-down list.
    - **IP address**—Enter the public IP address of the managed device.
    - **Network mask**—Enter the network mask or range.
16. Click **OK**.

    The new alias appears in the **Destination alias** drop-down list under **Configuration** > **Roles and Policies** > **Policies** tab.

17. To create the next rule:
    a. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Roles and Policies** > **Policies** tab.
    b. Click the policy created in step on page 875.
    c. Click **+** in the **Policy > <policy name> Rules** table.

       The **New Rule for <policy name>** window is displayed.
    d. Select **Access control** or **Application** in the **Rule Type**field.
    e. Click **OK**.
18. In the **New Policy > New Forwarding Rule** table, configure the following parameters:
    - **IP version**—Select **IPv4** or **IPv6** from the drop-down list.
    - **Source**—Select **Any** from the drop-down list.
    - **Destination**—Select **Alias** from the drop-down list.
    - **Destination alias**—Select the alias that you created from step on page 875.
    - **Action**—Select **Route Source NAT** from the drop-down list.
19. Click **Submit**.
20. In the **Managed Networks** node hierarchy, navigate to the **Configuration** > **Roles and Policies > Roles** tab.

    Roles can be created only in the managed device.

21. Click **+** to create a new role.
22. Enter the role name in the **Name** field.
23. Click **Submit**.
24. Click the new role created.
25. Click **Show Advanced View**.
26. Click **+**.
27. Select **Add existing session policy** option and select the policy created from the **Policy name** drop-down list.
28. Click **Submit**.
29. Click **Pending Changes**.

30. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following commands configure the session ACL allowing tunneling.

```
(host) [md] (config) #ap system-profile <profile>
lms-preemption
lms-hold-down period <seconds>netdestination <policy>
network <ipaddr> <netmask>

(host) [md] (config) #ip access-list session <policy>
any any svc-dhcp permit
any alias <name> any permit
user any any route src-nat

(host) [md] (config) #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 the following command.

```
(host) [md] (config) #ip access-list session <policy>
user alias <name> any redirect 0
user alias <name> any route
user alias <name> any route src-nat
```

Configuring an ACL to Restrict Local Debug Homepage Access

A user in split or bridge role using a Remote AP can log on to the local debug (LD) homepage (for example, (http://rapconsole.arubanetworks.com) and perform a reboot or reset operations. The LD homepage provides various information about the Remote AP and also has a button to reboot the Remote AP. You can now restrict a Remote AP user from resetting or rebooting a Remote AP by using the **localip** keyword in the in the user role ACL.



You will require the PEFNG license to use this feature. For complete information on the centralized licensing requirements, refer to the *Aruba Mobility Master Licensing Guide.*

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 Remote AP. The `localip` keyword identifies the set of all local IP addresses on the system to which the ACL is applied. The existing keywords Managed Device and `mswitch` indicate only the primary IP address on the Managed Device.



This release of ArubaOS provides localip keyword support only for Remote AP and not for Managed Device.

The following procedure describes how to configure an ACL to restrict local debug homepage access:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration** > **Roles and Policies** > **Policies** tab.
2. Click **+** to create a new policy.
3. Enter the policy name in the **Policy Name** field.
4. From the **Policy Type** drop-down list, select **Session**.
5. Click **Submit**.

6. To create the first rule:
   a. Select the policy created.
   b. Click **+**in the **Policies > <policy name>** table .
   c. Select the **Rule Type** in the **New Rule for <policy name>** dialog box.
   d. Click **OK**.
7. Enter the following details in the **Roles > <policy name> > <rule name>** table:
   a. From the **IP version** drop-down list, select **IPv4** or **IPv6**.
   b. Select **Any** from the **Source** drop-down list.
   c. Select **Any** from the **Destination** drop-down list.
   d. Select **Service** from the **Service/app** drop-down list.
   e. Select **svc-dhcp** from the **Service alias** drop-down list.
   f. Select **Permit** from the **Action** drop-down list.
   g. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure an ACL to restrict local debug homepage access:

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

All users have an ACL entry of type `any any deny` by default. This rule restricts access to all users. When the ACL is configured for a user role, if a `user any permit` ACL rule is configured, add a deny ACL before that for `localip` for restricting the user from accessing the LD homepage.

Example:

```
(host) [md] (config) #ip access-list session logon-control
user localip svc-http deny
user any permit
```

Configuring the AAA Profile for Tunneling

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

If you enable RADIUS accounting in the AAA profile, the Managed Device sends a RADIUS accounting start record to the RADIUS server when a user associates with the remote AP, and sends a stop record when the user logs out or is deleted from the user database. If you enable interim accounting, the Managed Device sends updates at regular intervals. Each interim record includes cumulative user statistics, including received bytes and packets counters. For more information on RADIUS accounting, see on page 215

The following procedure describes how to configure the AAA profile for tunneling:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Select **Wireless LAN >AAA** under **All Profiles**.

3. Click **+** in **AAA Profile** and enter the following details:
   a. Enter the **Profile name**.
   b. For **Initial role**, select the appropriate role (for example, logon) from the drop-down list.
   c. For **802.1X Authentication Default Role**, select the appropriate role (for example, "default") from the drop-down list.
   d. Click **Save**.
   e. Under the AAA profile that you created, locate **802.1X Authentication Server Group**, and select the Server Group to be used (for example "default") from the drop-down list.
   f. Click **Save**.
4. (Optional) To enable RADIUS accounting:
   a. Select the AAA profile from the profile list to display the list of authentication and accounting profiles associated with the AAA profile.
   b. Select the **Radius Accounting Server Group** profile associated with the AAA profile. Click the **Server Group** drop-down list to select a RADIUS server group. (For more information on configuring a RADIUS server or server group, see Configuring Authentication Servers on page 193.)
5. Click **Save**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configures the AAA profile for tunneling:

```
(host) [md] (config) #aaa profile <name>
authentication-dot1x <dot1x-profile>
dot1x-default-role <role>
dot1x-server-group <group>
radius-accounting <group>
radius-interim-accounting
```

Configuring the Virtual AP Profile

The following procedure describes how to configure the virtual AP profile:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles** tab.
2. Select **Wireless LAN > Virtual AP**under **All Profiles.**
3. To create a new virtual AP profile, click **+** in **Virtual AP profile**. Enter the name for the virtual AP profile, and click **Save**.



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. Select **Wireless LAN > SSID** under **All Profiles.**
   b. Click **+** in **SSID Profile** and enter the following details:
      i. Enter the name of the profile in the **Profile name** field.
      ii. To enable SSID, select the **SSID enable** check box.

iii. Select the appropriate check box under the **Encryption** field, to choose the network authentication and encryption method.

c. Click **Save**.

4. Select the new virtual AP name listed under **Wireless LAN > Virtual AP**, to view the configuration parameters.
5. In the **General** accordion under **Virtual AP profile: name**, execute the following:
   a. Ensure that the **Virtual AP enable** is selected.
   b. Enter the VLAN ID to be used for the Virtual AP profile in the **VLAN** field.
   c. Select **split-tunnel** from the **Forward mode** drop-down list.
   d. Click **Save**.
6. Under **All Profiles**, select **AP > AP system** profile.
7. Select the AP system profile that you want to edit.
   a. Under the **LMS Settings** accordion, enter the LMS IP address in the **LMS IP** field.
   b. Under the **Remote AP** accordion, click + under **Remote-AP DHCP DNS Server** and enter the Remote -AP DHCP DNS server in the **Remote-AP DHCP DNS Server** field.
   c. Click **OK**.
   d. Click **Save**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the virtual AP profile:

```
(host) [md] (config) #wlan ssid-profile <profile>
essid <name>
opmode <method>

(host) [md] (config) #wlan virtual-ap <profile>
ssid-profile <name>
forward-mode <mode>

(host) [md] (config) # vlan <vlan id>
aaa-profile <profile>

(host) [md] (config) #ap-group <name>
virtual-ap <profile>
```

or

```
(host) [md] (config) #ap-name <name>
virtual-ap <profile>
```

Defining Corporate DNS Servers

Clients send DNS requests to the corporate DNS server address that it learned from DHCP. If configured for split tunneling, corporate domains and traffic destined for corporate use the corporate DNS server. For non-corporate domains and local traffic, other DNS servers can be used.

The following procedure describes how to define corporate DNS servers:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles** tab.
2. Select **AP > AP system** under **All Profiles.**
3. Select an AP System Profile.
4. Click **+** under **Corporate DNS Domain** table and enter the **Corporate DNS Domain**.
5. Click **OK**. The DNS name appears in **Corporate DNS Domain** table. You can add multiple names the same way.
6. Click **Save**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands define corporate DNS servers:

```
(host) [md] (config) #ap system-profile <profile>
dns-domain <domain name>
```

Understanding Bridge

The bridge feature allows you to route the traffic flow only to the internet and not to the corporate network. Only the 802.1X authentication request is sent to the corporate network. This feature is useful for guest users.

NOTE

ArubaOS does not support Wired 802.1X authentication in bridge mode for a s. 802.1X authentication is supported only in tunnel and split modes.

Figure 105 Sample Bridge Environment

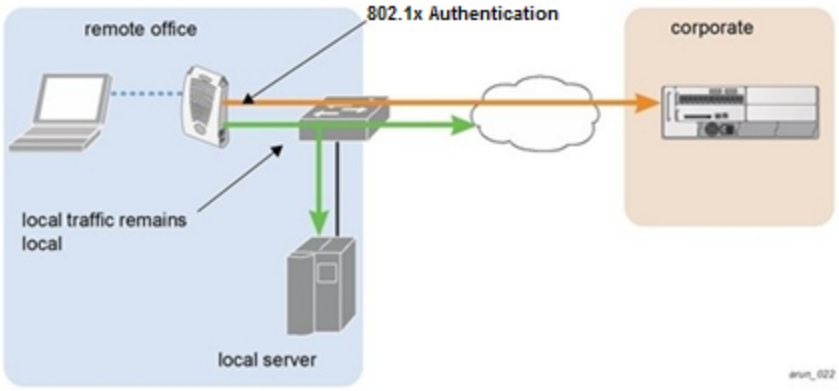


Figure 105 displays the local traffic being routed to the internet and the 802.1X authentication request sent to the corporate network.

Configuring Bridge

The following procedure describes how to configure a bridge. Each step is described in detail later in this chapter.

NOTE

The bridge feature requires the PEFNG license. If you do not have the PEFNG license on your managed device, you must install it before you configure bridge. For details on installing licenses, refer to the *Aruba Mobility Master Licensing Guide.*

1. Define a session ACL that routes the traffic.
   a. Create rules to permit DHCP and local data traffic.
   b. Apply the session ACL to a user role. For information about user roles and policies, see Roles and Policies on page 491.
2. Configure the remote AP's AAA profile. Specify the authentication method (**802.1X** or **PSK**) and the default user role for authenticated users. The user role specified in the AAA profile must contain the session ACL defined in the previous step. Optionally, use the remote AP's AAA profile to enable RADIUS accounting.
3. Configure the virtual AP profile:
   a. Specify the AP group or ap-name to which the virtual AP profile applies.
   b. Set the VLAN in the virtual AP.
   c. When specifying the use of a bridge configuration, use bridge forward mode.
   d. Create and apply the applicable SSID profile. Optionally under AP system profile, configure the Remote AP DHCP pool. Remote AP DHCP VLAN must be same as virtual AP's VLAN. If the client needs to obtain from the Remote AP DHCP Server.



When creating a new virtual AP profile In the WebUI, you can simultaneously configure the SSID. For information about AP profiles, see AP Configuration Profiles on page 639.

Configuring the Session ACL

First you need to configure a session ACL that "permits" corporate traffic to be forwarded to the managed device and that routes, or locally bridges, local traffic.

The following procedure describes how to configure session ACL:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles and Policies > Policies** tab.
2. Click **+** to create a new policy.
3. Enter the name in the **Policy name** field.
4. Select **Session** from the **Policy type** drop-down list.
5. Click **Submit**.
6. Select the policy created and click **+** under **Policies<policy name>** table.
7. Select **Access Control** option in the **Rule Type** field.
8. Click **OK**.
9. To complete creating the rule:
   a. Select **IPv4** or **IPv6** from the **IP version** drop-down list.
   b. Select **Any** from the **Source** drop-down list.
   c. Select **Any** from the **Destination** drop-down list.
   d. Select **Service** from the **Service/app** drop-down list.
   e. Select **svc-dhcp** from the **Service alias** drop-down list.
   f. Select **Permit** for IPv4 or **Captive** for IPv6 from the **Action** drop-down list.
   g. Click **Submit**.

10. To create a new forwarding rule:
    a. Select policy created and click **+** in the **Policies <policy name>** table.
    b. Select **Access Control** option in the **Rule Type** field.
    c. Click **OK**.
    d. Select **IPv4** or **IPv6** from the **IP version** drop-down list.
    e. Select **any** from the **Source** drop-down list.
    f. Select **alias** from the **Destination** drop-down list.
    g. Click **+** in the **Destinationalias** drop-down list.
    h. In the **Add New Destination** window, click **+** in the **Rule** table.
    i. Select **Network** from the **Rule type** drop-down list.
    j. Enter the public IP address of the managed device in the **IP address** field.
    k. Enter the netmask or range in the **Network mask** field.
    l. Click **OK**. The new alias appears in the **Destination alias** drop-down list.
    m. Click **Submit**.
11. Navigate to the **Configuration** >**Roles and Policies** >**Roles** tab.



Roles can be created only in the managed device.

    a. Click **+** to create a new role.
    b. Enter the role name in the **Name** field.
    c. Click **Submit**.
    d. Click the new role created.
    e. Click **Show Advanced View**.
    f. Click **+**.
    g. Select **Add an existing policy** option and select the policy created from the **Policy name** drop-down list.
    h. Click **Submit**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure session ACL:

If `dhcp server` in `ap system profile` is enabled:

```
(host) [md] (config) #ip access-list session <policy> any any svc-dhcp permit
(host) [md] (config) #user any any route src-nat
```

If `dhcp server` in `ap system profile` is disabled:

```
(host) [md] (config) #ip access-list session <policy>
(host) [md] (config) #any any any permit
(host) [md] (config) #user-role <role>
(host) [md] (config) #session-acl <policy>
```



To configure an ACL to Restrict Local Debug Homepage Access, see Configuring an ACL to Restrict Local Debug Homepage Access on page 1.

Configuring the AAA Profile for Bridge

After you configure the session ACL, define the AAA profile used for bridge. When defining the AAA parameters, specify the previously configured user role that contains the session ACL used for bridge.

If you enable RADIUS accounting in the AAA profile, the Mobility Master sends a RADIUS accounting start record to the RADIUS server when a user associates with the remote AP, and sends a stop record when the user logs out or is deleted from the user database. If you enable interim accounting, the Mobility Master sends updates at regular intervals. Each interim record includes cumulative user statistics, including received bytes and packets counters. For more information on RADIUS accounting, see on page 215.

The following procedure describes how to configure the AAA profile for bridge:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Select **Wireless LAN > AAA** under **All Profiles.**
3. Click **+** in **AAA Profile**.
4. Enter the **Profile name**.
5. Select the appropriate role (for example, “logon”) from the **Initial role** drop-down list.
6. Select the user role you previously configured for split tunneling or bridge from the **802.1X Authentication Default Role** drop-down list.
7. Click **Save**.
8. Select the AAA profile created and locate the **802.1X Authentication Server Group**, and select the **Server Group** to be used (for example “default”) from the drop-down list.
9. Click **Save**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configures the AAA profile for bridge:

```
(host) [mynode] (config) #aaa profile <name>
(host) [mynode] (config) #authentication-dot1x <dot1x-profile>
(host) [mynode] (config) #dot1x-default-role <role>
(host) [mynode] (config) #dot1x-server-group <group>
(host) [mynode] (config) #radius-accounting <group>
(host) [mynode] (config) #radius-interim-accounting
```

Configuring the Virtual AP Profile

The following procedure describes how to configure the virtual AP profile:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles** tab.
2. Select **Wireless LAN > Virtual AP**under **All Profiles.**
3. To create a new virtual AP profile, click **+** in **Virtual AP profile**. Enter the name for the virtual AP profile, and click **Save**.



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. Select **Wireless LAN > SSID** under **All Profiles.**
   b. Click **+** in **SSID Profile** and enter the following details:
      i. Enter the name of the profile in the **Profile name** field.
      ii. To enable SSID, select the **SSID enable** check box.
      iii. Select the appropriate check box under the **Encryption** field, to choose the network authentication and encryption method.
   c. Click **Save**.
4. Select the new virtual AP name listed under **Wireless LAN > Virtual AP**, to view the configuration parameters.
5. In the **General** accordion under **Virtual AP profile: name**, execute the following:
   a. Ensure that the **Virtual AP enable** is selected.
   b. Enter the VLAN ID to be used for the Virtual AP profile in the **VLAN** field.
   c. Select **split-tunnel** from the **Forward mode** drop-down list.
   d. Click **Save**.
6. Under **All Profiles**, select **AP > AP system** profile.
7. Select the AP system profile that you want to edit.
   a. Under the **LMS Settings** accordion, enter the LMS IP address in the **LMS IP** field.
   b. Under the **Remote AP** accordion, click + under **Remote-AP DHCP DNS Server** and enter the Remote -AP DHCP DNS server in the **Remote-AP DHCP DNS Server** field.
   c. Click **OK**.
   d. Click **Save**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the virtual AP profile:

```
(host) [md] (config) #wlan ssid-profile <profile>
essid <name>
opmode <method>

(host) [md] (config) #wlan virtual-ap <profile>
ssid-profile <name>
forward-mode <mode>

(host) [md] (config) # vlan <vlan id>
aaa-profile <profile>

(host) [md] (config) #ap-group <name>
virtual-ap <profile>
```

or

```
(host) [md] (config) #ap-name <name>
virtual-ap <profile>
```

Provisioning Wi-Fi Multimedia

WMM is a WFA specification based on the IEEE 802.11e wireless 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 and DSCP tags. Remote APs support WMM.

WMM supports four access categories: 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 or mappings. If this happens, your traffic may not be prioritized correctly.

Reserving Uplink Bandwidth

You can reserve and prioritize uplink bandwidth traffic to provide higher QoS for specific applications, traffic, or ports. This is done by applying bandwidth reservation on existing session ACLs. Typically, the bandwidth reservation is applied for uplink voice traffic.

Note the following before you configure bandwidth reservation:

- You must know the total bandwidth available.
- Bandwidth reservation is applicable only on session ACLs.
- Bandwidth reservation on voice traffic ACLs receives higher priority over other reserved traffic.
- You can configure up to three unique priority for bandwidth reservation.
- The bandwidth reservation must be specified in absolute value (Kbps).
- Priorities for bandwidth reservation are optional, and bandwidth reservations without priorities are treated equal.

Understanding Bandwidth Reservation for Uplink Voice Traffic

Voice ACLs are applicable on the voice signaling traffic used to establish a 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 signaling traffic and ports used for voice data traffic. This mechanism filters traffic as per the security requirements.

Configuring Bandwidth Reservation

You can configure bandwidth reservation ACLs using the WebUI or the CLI.

The following procedure describes how to configure bandwidth reservation:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Click **AP** and select **AP system**.
3. Click **Remote AP** accordion. You can create a new AP system profile to configure bandwidth reservation or edit an existing AP system profile. In **Remote-AP bw reservation 1**, **Remote-AP bw reservation 2** and **Remote-AP bw reservation 3** fields, specify bandwidth reservation values.

4. Click **Save**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure bandwidth reservation

```
(host) [mynode] (config)#ap system-profile remotebw
(host) [mynode] (AP system profile "remotebw") #rap-bw-total 1024
(host) [mynode] (AP system profile "remotebw") #rap-bw-resv-1 acl voice 128 priority 1
```

To view bandwidth reservations:

```
(host) [mynode] #show datapath rap-bw-resv ap-name remote-ap-1
```

Provisioning 4G USB Modems on Remote Access Points

ArubaOS provides support for 4G networks by allowing you to provision 4G USB modems on the Remote AP. You can also provision the Remote AP to support both 4G and 3G USB modems. This enables the Remote AP to choose the available network automatically. 4G takes precedence over 3G when the Remote AP tries to auto select the network. You can also configure the Remote AP to work exclusively on a 3G or 4G network. It is recommended that you provision the USB modems for the Remote AP based on your network requirements.

4G USB Modem Provisioning Best Practices and Exceptions

- Remote AP does not support dynamic plug-and-play for the 4G USB modems. You must provision a Remote AP with the 4G USB parameters on the managed device manually based on its type and family (4G-WiMAX or 4G-LTE).
- When a Remote AP connects to a 4G network, it appears as a Remote AP (R) and Cellular (C) on the managed device.
- For a 3G or 4G network switch, using the UML290 modem with the firmware version L0290VWB522F.242 or later is recommended. Using a lower version of the firmware auto-selects the network mode based on the network availability. The latest version allows the Remote AP to lock the modem in a particular network mode (for example, 3G only).



The 4G-WiMAX family of modems do not support the 3G-4G network switch-over functionality.

A new method of provisioning multimode USB modems (such as a Verizon UML290, Verizon MC551L, AT&T 313u, Huawei K5150, AT &T ZTE MF861 and Inseego U730L) for a Remote AP has been introduced. These changes simplify modem provisioning for both 3G and 4G networks. Earlier the modem configuration procedure required that you define a driver for a 3G modem in the USB modem field under the AP provisioning profile, or define a driver for a 4G modem in the 4G USB type field. You can now configure drivers for both a 3G or a 4G modem using the USB field, and the 4G USB Type field is deprecated. The managed device can auto configure the USB modem when it is plugged into the associated Remote AP. Since most 4G- LTE modem support dynamic network-switching between 4G and 3G, by default (for zero touch) Remote AP is configured in 3G/4G mode. In such cases, the Remote AP selects the best available cellular network coverage in that specific region.

Provisioning Remote AP for USB Modems

To enable 3G or 4G network support, you must provision the Remote AP with the USB parameters on the managed device. You can use the WebUI or CLI to provision the USB parameters.

The following procedure describes how to provision Remote AP for USB modems:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Access Points > Remote APs** tab.
2. Select the remote AP and click **Provision**.
3. Select **Uplink** tab. This tab is displayed only when a Remote AP is selected.
4. Select a profile from the **USB Profile** drop-down list. This field is displayed only when the device is USB enabled.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Table 158: *Cellular Network Preference Parameters*

Parameter	Description
auto (default)	In this mode, the modem firmware will control the cellular network service selection; so the cellular network service failover and fallback is not interrupted by the Remote AP.
3g_only	Locks the modem to operate only in 3G.
4g_only	Locks the modem to operate only in 4G.
advanced	The Remote AP controls the cellular network service selection based on an RSSI threshold-based approach. ▪ Initially the modem is set to the default auto mode. This allows the modem firmware to select the available network. ▪ The Remote AP determines the RSSI value for the available network type (for example 4G), checks whether the RSSI is within required range, and if so, connects to that network. ▪ If the RSSI for the modem's selected network is not within the required range, the Remote AP will then check the RSSI limit of an alternate network (for example, 3G), and reconnect to that alternate network. The Remote AP will repeat the above steps each time it tries to connect using a 4G multimode modem in this mode.

The following CLI commands provision Remote AP for USB modems:

To enable 4G-exclusive network support on the Remote AP, execute the following commands:

```
(host) [md] (config) #ap provisioning-profile <profile-name>
(host) [md] (Provisioning profile "<profile-name>") usb-type <USB modem type>
(host) [md] (Provisioning profile "<profile-name>") #usb-type none
(host) [md] (Provisioning profile "<profile-name>") #cellular_nw_preference 4g_only
```

To enable 3G-exclusive network support on the Remote AP, execute the following commands:

```
(host) [md] (config) #ap provisioning-profile <profile-name>
(host) [md] (Provisioning profile "<profile-name>") usb-type <USB modem type>
(host) [md] (Provisioning profile "<profile-name>") #usb-type none
(host) [md] (Provisioning profile "<profile-name>") #cellular_nw_preference 3g_only
```

To enable 3G or 4G network switch support, execute the following commands:

```
(host) [md] (config) #ap provisioning-profile <profile-name>
```

```
(host) [md] (Provisioning profile "<profile-name>") usb-type <USB modem type>
(host) [md] (Provisioning profile "<profile-name>") #usb-type none
(host) [md] (Provisioning profile "<profile-name>") #cellular_nw_preference auto
```

Remote AP 3G or 4G Backhaul Link Quality Monitoring

The Remote AP is enhanced to support link monitoring on 2G, 3G, and 4G modems to provide information about the state of the USB modem and cellular network.

The USB modem has the following four states:

- **Active** - The USB modem is used as the primary path for connecting VPN to the managed device
- **Standby** or **Backup** - The network is available but the USB modem is not used for connecting VPN to the managed device
- **Error** - The USB modem is available but the modem is faulty
- **Not Plugged** - The USB modem is unavailable

To view the USB modem details on the Remote AP, execute the following command:

```
(host) [md] #show ap debug usb ap-name <ap-name>
```

Provisioning Remote AP at Home

The following section provides information on provisioning your Remote Ap at home using a static IP address, PPPoE connection, or USB modem.

Prerequisites

Follow the steps below to acquire a static IP address before provisioning the Remote AP at home:

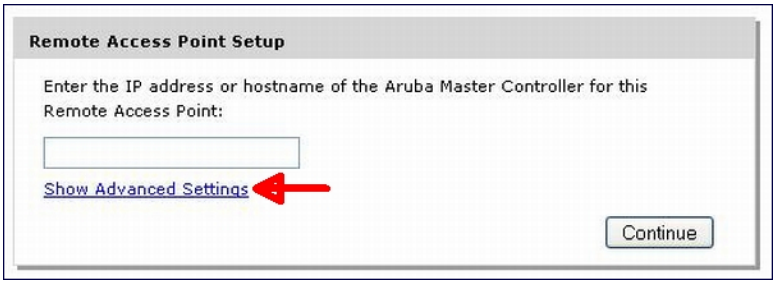
1. Connect the Remote AP at the site of deployment and ensure that it has connectivity to the Internet to reach the managed device.
2. Connect a laptop to Port 1 of the Remote AP to get an IP address from the Remote AP's internal DHCP pool.

Provisioning Remote AP Using ZTP

You provision the Remote AP using provisioning wizard:

1. Navigate to the Remote AP configuration URL: http://rapconsole.arubanetworks.com.
2. Enter the IP address or hostname of the managed device.
3. Click the **Show Advanced Settings** link, shown in Figure 106.

Figure 106 *Show Advanced Settings*



4. In the **Advanced Settings** wizard, you can select one of the following:
   a. **Static IP**—Select this tab to provision your Remote AP using a static IP address.
   b. **PPPoE**—Select this tab to provision your Remote AP on a PPPoE connection.
   c. **USB**—Select this tab to provision your Remote AP using 3G/EVDO USB modem.

Provisioning the Remote AP using a Static IP Address

Select the **Static IP** tab and enter the required details. See Table 159 for information on parameters.

Figure 107 *Provision Remote AP using Static IP*

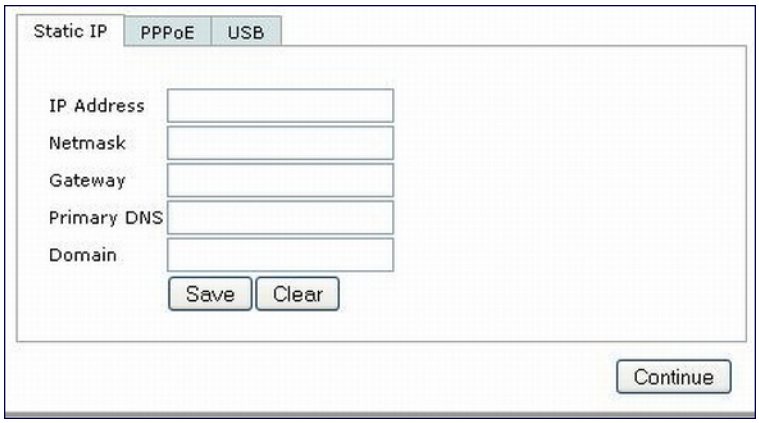


Table 159: *Provision using Static IP*

Parameter	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 **Save** after you have entered all the details.

Provision the Remote AP on a PPPoE Connection

Select the **PPPoE** tab and enter the required details. See Table 160 for information on parameters.

Figure 108 *Provision Remote AP on a PPPoE Connection*

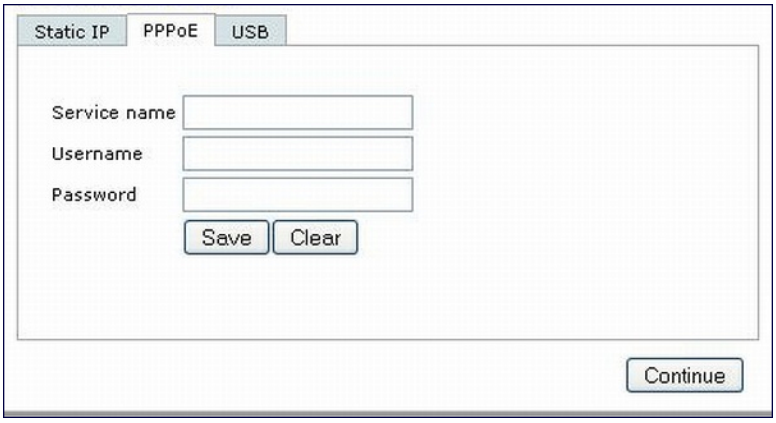


Table 160: *Provision using PPPoE Connection*

Parameter	Description
Service Name	Enter the PPPoE service name provided to you by your service provider. This parameter is optional.
Username	Enter the user name for the PPPoE connection.
Password	Enter your PPPoE password.

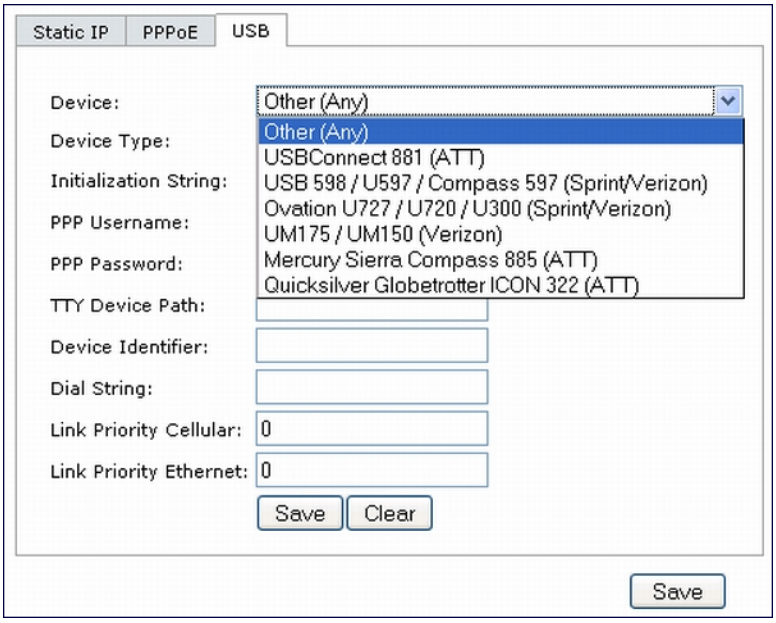
Click **Save** after you have entered all the details.

Using 3G/EVDO USB Modems

The following procedure illustrates provisioning your Remote AP using a 3G/EVDO USB modem.

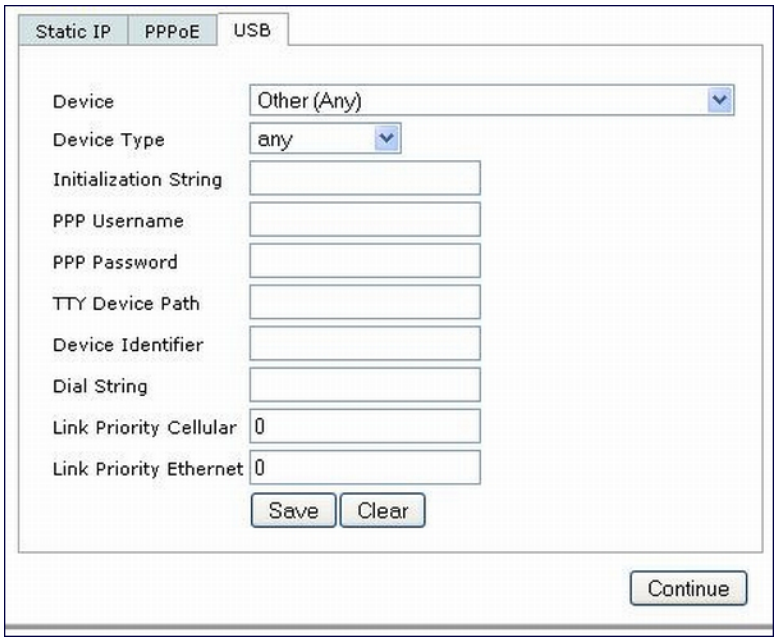
1. Select the **USB** tab and select your modem from the drop-down list. Configuration details automatically appear for some common modems.

   Figure 109 *Provision using a preconfigured USB Modem*



2. If your modem name is not listed, select **Other** and manually enter the following details. These are available from the manufacturer of your modem or from your IT administrator:

Figure 110 *Provision using a USB Modem with Custom Settings*



- Device Type
- Initializing String
- PPP Username
- PPP Password
- TTY Device Path
- Device Identifier
- Dial String
- Link Priority Cellular—This is a number that identifies the priority of the connection. If the *Link Priority Cellular* has a higher number than *Link Priority Ethernet*, then cellular connection is used.
- Link Priority Ethernet—This is a number that identifies the priority of the connection. If the *Link Priority Ethernet* has a higher number than *Link Priority Cellular*, then Ethernet connection is used.

3. Click **Save** after you have entered all the details and click **Continue** to complete provisioning of your Remote AP.

Converting an Instant AP to Remote AP or Campus AP

For Instant AP to Remote AP or Campus AP conversion, the virtual controller sends the convert command to all the other Instant APs. The virtual controller along with the other slave Instant APs then set up a VPN tunnel to the remote controller, and download the firmware by FTP. The virtual controller uses IPsec to communicate to the controller over the Internet.

NOTE

A mesh point cannot be converted to a Remote AP because mesh does not support VPN connection.

Important

- Converting an AP to Instant AP is only supported on UAP models.
- Converting non-UAP models is not supported on the Web UI and can only be done through CLI.

Converting Instant AP to Remote AP

To convert an Instant AP to Remote AP, follow the instructions below:

1. Navigate to the **Maintenance** tab in the top right corner of the Instant UI.
2. Click the **Convert** tab.
3. Select **Remote APs managed by a Mobility Controller** from the drop-down list.
4. Enter the hostname FQDN or the IP address of the managed device in the **Hostname or IP Address of Mobility Controller** text box. This information is provided by your network administrator.



Ensure the controller IP Address is reachable by the IAPs.

5. Click **Convert Now** to complete the conversion.
6. The Instant AP reboots and begins operating in Remote AP mode.
7. After conversion, the Instant AP is managed by the Aruba controller which has been specified in the Instant UI.



In order for the Remote AP conversion to work, ensure that you configure the Instant AP in the Remote AP white-list and enable the FTP service on the controller.

If the VPN setup fails and an error message pops up, please click OK, copy the error logs and share them with your Aruba support engineer.

Converting an Instant AP to Campus AP

To convert an Instant AP to a Campus AP, do the following:

1. Navigate to the **Maintenance** tab in the top right corner of the Instant UI.
2. Click the **Convert** tab.
3. Select **Campus APs managed by a Mobility Controller** from the drop-down list.
4. Enter the hostname FQDN or the IP address of the controller in the **Hostname or IP Address of Mobility Controller** text box. This is provided by your network administrator.



Ensure that the controller IP Address is reachable by the APs.

5. Click **Convert Now** to complete the conversion.

Enabling Bandwidth Contract Support for Remote APs

Bandwidth Contract support on Remote APs is achieved by extending the Bandwidth Contract support on split-tunnel and bridge modes. You can apply Bandwidth Contract for a Remote AP on a per-user or per-role basis. Bandwidth Contract is applied on a per-role basis by default. This implies that all the users belonging to the same role will share the bandwidth pool. When Bandwidth Contract configured on the managed device is attached to a user-role, it automatically gets pushed to the Remote APs terminating on it.

The following show commands have been enhanced to retrieve the Bandwidth Contract information from the Remote AP:

```
(host) [md] #show datapath user ap-name <ap-name>
(host) [md] #show datapath bwm ap-name <ap-name>
```

Configuring Bandwidth Contracts for Remote AP

The following examples illustrate how to configure, apply, and verify the Bandwidth Contracts on the RAPs.

Defining Bandwidth Contracts

Use the following command to define a 256 Kbps contract:

```
(host) [md] (config) #aaa bandwidth-contract 256k kbits 256
```

Use the following command to define a 512 Kbps contract

```
(host) [md] (config) #aaa bandwidth-contract 512k kbits 512
```

Applying Contracts

You can apply the contract on a per-role or per-user basis.

Use the following commands to apply the contracts on a per-role basis for upstream and downstream:

For upstream contract of 512 Kbps:

```
(host) [md] (config) #user-role authenticated bw-contract 512k upstream
```

For downstream contract of 256 Kbps:

```
(host) [md] (config) #user-role authenticated bw-contract 256k downstream
```

Use the following commands to apply the contracts on a per-user basis for upstream and downstream:

For upstream contract of 512 Kbps:

```
(host) [md] (config) #user-role authenticated bw-contract 512k per-user upstream
```

For downstream contract of 256 Kbps:

```
(host) [md] (config) #user-role authenticated bw-contract 256k per-user downstream
```

Verifying Contracts on AP

The following example displays the bandwidth contracts on an AP for per-role configuration:

```
(host) [md] #show datapath bwm ap-name rap5-2
Datapath Bandwidth Management Table Entries
-------------------------------------------
Contract Types :
0 - CP Dos 1 - Configured contracts 2 - Internal contracts
-----------------------------------------------
Flags: Q - No drop, P - No shape(Only Policed),
T - Auto tuned
---  --------  ---------  ----------  -----------  -----------------
Cont                      Avail   Queued/Pkts
Type  Id   Bits/sec  Policed    Bytes    Bytes      Flags
```

```
---- ---- --------- ---------- ------- ------------ -----
1    1      512000          0   16000       0/0    P
1    2      256000          0    8000       0/0    P
```

The following example displays the bandwidth contracts on AP for per-user configuration (contract IDs 3 and 4 are per-user contracts):

```
(host) [md] #show datapath bwm ap-name rap5-2
Datapath Bandwidth Management Table Entries
-------------------------------------------
Contract Types :
0 - CP Dos 1 - Configured contracts 2 - Internal contracts
-------------------------------------------
Flags: Q - No drop, P - No shape(Only Policed),
T - Auto tuned
---  --------  ---------  ----------  -----------  -----------------
Cont                       Avail   Queued/Pkts
Type  Id   Bits/sec  Policed    Bytes    Bytes      Flags
---- ---- --------- ---------- ------- ------------ -----
1    1       512000        300   16000       0/0    P
1    2       256000        277    8000       0/0    P
1    3       512000          0   16000       0/0    P
1    4       256000          0    8000       0/0    P
```

Verifying Contracts Applied to Users

You can verify if the contracts are applied to the user after the user connects to the AP using CLI.

The following is a sample output for a per-role configuration:

```
(host) [md] #show datapath user ap-name rap5-2
Datapath User Table Entries
---------------------------
Flags: P - Permanent, W - WEP, T- TKIP, A - AESCCM, G - AESGCM, V - ProxyArp to/for MN
(Visitor),
N - VPN, L - local, Y - Any IP user, R - Routed user, M - Media Capable,
S - Src NAT with VLAN IP, E - L2 Enforced, F - IPIP Force Delete, O - VOIP user
FM(Forward Mode): S - Split, B - Bridge, N - N/A
IP              MAC           ACLs    Contract   Location  Age  Sessions   Flags
Vlan   FM
---------------  -----------------  -------  ---------  --------  ---  ---------  ----
-   ----   --
10.15.72.50      00:0B:86:61:12:AC  2703/0      0/0    0         16       1/65535
P      0   N
10.15.72.253     00:18:8B:A9:A8:DF    52/0      1/2    0         1        0/65535
      1   S
192.168.11.1     00:0B:86:66:03:3F  2700/0      0/0    0         20024    0/65535
P    177   N
10.15.196.249    00:0B:86:66:03:3F  2700/0      0/0    0         3        1/65535
P      1   N
```

The following is a sample output for a per-user configuration:

```
(host) [mynode] #show datapath user ap-name rap5-2
Datapath User Table Entries
---------------------------
Flags: P - Permanent, W - WEP, T- TKIP, A - AESCCM, G - AESGCM, V - ProxyArp to/for MN
(Visitor),
N - VPN, L - local, Y - Any IP user, R - Routed user, M - Media Capable,
S - Src NAT with VLAN IP, E - L2 Enforced, F - IPIP Force Delete, O - VOIP user
FM(Forward Mode): S - Split, B - Bridge, N - N/A
IP              MAC           ACLs    Contract   Location  Age  Sessions   Flags
Vlan   FM
```

```
--------------  -----------------  -------  ---------  --------  ---  ---------  ----
-   ----   --
10.15.72.50      00:0B:86:61:12:AC  2703/0      0/0    0         11       0/65535  P
      0    N
10.15.72.253     00:18:8B:A9:A8:DF    52/0      3/4    0         46       0/65535
      1    S
192.168.11.1     00:0B:86:66:03:3F  2700/0      0/0    0         20883    0/65535  P
    177    N
10.15.196.249    00:0B:86:66:03:3F  2700/0      0/0    0         15       1/65535  P
      1    N
```

Verifying Bandwidth Contracts During Data Transfer

You can verify the Bandwidth Contracts that are in use during data transfer using CLI.

The following is a sample output for a per-role configuration:

```
(host) [md] #show datapath session ap-name rap5-2 table 10.15.72.99
Datapath Session Table Entries
------------------------------
Flags: F - fast age, S - src NAT, N - dest NAT
D - deny, R - redirect, Y - no syn
H - high prio, P - set prio, T - set ToS
C - client, M - mirror, V - VOIP
Q - Real-Time Quality analysis
I - Deep inspect, U - Locally destined
E - Media Deep Inspect, G - media signal
RAP Flags: 1 - Class 1, 2 - Class 2, 3 - Class 3
Source IP     Destination IP  Prot SPort DPort  Cntr Prio ToS Age Destination TAge
Flags
--------------  --------------  ---- ----- -----  ---- ---- --- --- ----------- ---- -
----
10.15.72.253    10.15.72.99     6    5001  36092  1/1     0 0   0   dev12       6
10.15.72.253    10.15.72.99     6    3488  5001   1/1     0 0   0   dev5        6    C
10.15.72.99     10.15.72.253    6    5001  3488   1/2     0 0   0   dev5        6
10.15.72.99     10.15.72.253    6    36092 5001   1/2     0 0   0   dev12       6    C
```

The following is a sample output for a per-user configuration:

```
(host) [md] #show datapath session ap-name rap5-2 table 10.15.72.99

Datapath Session Table Entries
------------------------------
Flags: F - fast age, S - src NAT, N - dest NAT
D - deny, R - redirect, Y - no syn
H - high prio, P - set prio, T - set ToS
C - client, M - mirror, V - VOIP
Q - Real-Time Quality analysis
I - Deep inspect, U - Locally destined
E - Media Deep Inspect, G - media signal
RAP Flags: 1 - Class 1, 2 - Class 2, 3 - Class 3
Source IP     Destination IP  Prot SPort DPort  Cntr Prio ToS Age Destination TAge
Flags
--------------  --------------  ---- ----- -----  ---- ---- --- --- ----------- ---- -
----
10.15.72.253    10.15.72.99     6    3489  5001   1/3     0 0   0   dev5        37
FC
10.15.72.99     10.15.72.253    6    5001  3489   1/4     0 0   0   dev5        37   F
10.15.72.99     10.15.72.253    6    36096 5001   1/4     0 0   0   dev12       37   C
10.15.72.253    10.15.72.99     6    5001  36096  1/3     0 0   0   dev12       37
```

VIA is part of the Aruba remote networks solution intended for teleworkers and mobile users. VIA detects the network environment (trusted and untrusted) of the user and connects the users to the enterprise network. Trusted networks refers to a protected office network that allows users to directly access the corporate intranet. Untrusted networks are public Wi-Fi hotspots such as airports, cafes, or home network.

The VIA solution includes the VIA client, Mobility Master with managed device configuration.

- VIA client—Remote workers and mobile users can install VIA on their computers and smart devices (iOS and Android) to connect to their enterprise network from remote locations.
- Mobility Master and managed device configuration—To set up VIA for remote users, configure the VPN for VIA in the Mobility Master and configure the authentication profile and connection profile in the managed network.

  VIA configuration settings are in the following sections of the WebUI:
  - In the **Mobility Master** node hierarchy, navigate to **Configuration > Services > VPN > VIA**.
  - In the **Managed Network** node hierarchy, navigate to **Configuration > Authentication > L3 Authentication:**
    - **VIA Authentication**
    - **VIA Connection**
    - **VIA Web Authentication**



For information on configuring the settings in these profiles, refer to the VIA 3.x User Guide.

ArubaOS now supports IPv6 connectivity for VIA that allows you to configure IPv6 address of the managed device in VIA connection profile. Hence, you can use either IPv4 or IPv6 address of the managed device to establish connection with the remote server. However, you must still use inner IPv4 address in VIA connection profile.

The following CLI commands configure IPv6 address of the managed device:

```
(host) [mynode] (config) #aaa authentication via connection-profile test
(host) [mynode] (VIA Connection Profile "test") #server addr <ipv6-addr> internal-ip <ip-addr> desc <description> position <number> 1
```

This section also includes the following topics:

License Requirements

Managed devices running ArubaOS 8.x require one of two available license types to support VIA users, the **PEFV** license, or the **VIA** license.

The **PEFV** license allows a network administrator to apply firewall policies to clients using a VPN to connect to the managed device. This PEFV license is purchased as a single device-specific license that enabled the functionality up to the full user capacity of the managed device.

ArubaOS 8.2.0.0 and later supports a sharable **VIA** license. Each VIA client or 3rd party VPN client consumes a single VIA license. (VIA licenses are not consumed by site-to-site VPNs.) If a standalone

controller or a managed device managed by Mobility Master has a PEFV license, that device will not consume VIA licenses from a licensing pool, as a single PEFV license supports all VIA and 3rd party VPN clients, up to the full user capacity for that controller or managed device.

Marking Outgoing Packets with ToS Bits

Starting from ArubaOS 8.3.0.0, you can configure the type of service-differentiated service code point (ToS-DSCP) for managed devices. This provides the ability for VIA to mark outgoing IKE and ESP packets with custom DSCP. When a VIA client downloads the connection-profile, this value also gets pushed. VIA sets the configured DSCP value to outer IP header's ToS byte. You can use this to mark IPsec packets with higher QoS/DSCP than Best Effort.

The following procedure describes how to configure the **tos-dscp** parameter in the WebUI:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles** and expand the **Other Profiles** menu.
2. Expand the **VIA Connection** profile option and select the name of an existing profile or click **Add** to create a new profile.
3. Click the **default** profile or other saved profile where you want to make changes.
4. In the **VIA Connection Profile:<profile-name>** pane on the right, enter a value for **tos-dscp**. The allowed value range is 0-63.
5. Click **Submit**.
6. Select **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the **tos-dscp** parameter in the managed device node:

```
(host) [mynode] (config) #aaa authentication via connection-profile <profile-name>
(host) [mynode] (VIA Connection Profile "<profile-name>") #tos-dscp <0-63>
```

For more details on configuring, installing, and using VIA, refer to the latest version of the *Aruba VIA for Mobility Master User Guide*.

VIA Client Audit

Starting from ArubaOS 8.4.0.0, when a user authenticates and accesses the VIA client, a notification with details about the last successful logon date and time stamp is provided.

The following CLI commands enables to view the username and the last login information:

```
(host) [mm] #show via-lastlogin
```

VIA VPN Client Visibility

Starting from ArubaOS 8.4.0.0, the VIA client users are separately displayed on the WebUI for VPN client visibility. You can view the client users in the **Dashboard > Clients > Remote Clients** page in the WebUI.

Previously, you could view the VIA VPN users using the CLI commands, **show user** and **show user-table**. However, now the VIA VPN users information is published to a GSM channel, user and can be seen using the CLI command, **show gsm debug channel user**.

VIA VPN Client Capability

Starting from ArubaOS 8.4.0.0, the VIA client provides a new option (VIA connection profile knob) to enable forwarding of Layer-2 GRE tunnel. This feature allows the VIA client to send GRE packets containing Ethernet frame by using the IPsec tunnel established with the managed device.

The following CLI commands enable the Layer-2 forwarding option in VIA connection profile:

```
(host) [mynode] (config) # aaa authentication via connection-profile default
(host) [mynode] (VIA Connection Profile "default") # l2-forwarding
```

VIA Unique Identifier

Starting ArubaOS 8.4.0.0, VIA uses the MAC address of a client as the calling station id when sending an authentication request to ClearPass Policy Manager. In earlier versions, the IP address of the client was used as the calling station id.

VIA VPN Client Authentication

Starting from ArubaOS 8.5.0.0, the VIA connection profile includes EAP-GTC authentication option. This option ensures that the VIA client establishes IKEv2 tunnel with the managed device.

The following procedure describes how to configure EAP-GTC in the WebUI:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles** tab.
2. In the **All Profiles** list, expand the **Other Profiles** menu.
3. Expand the **VIA Connection** profile option and select the name of an existing profile or click **+** to create a new profile.
4. In the **VIA Connection Profile:<profile-name>** pane on the right, select **eap-gtc** from the **IKEv2 Authentication method** field drop-down list.
5. Click **Submit**.
6. Select **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands set the EAP-GTC as the authentication method:

```
(host) [mynode] (config) # aaa authentication via connection-profile <profile_name>
(host) [mynode]  (VIA Connection Profile "profile_name") #ikev2auth eap-gtc
```

Chapter 36
Spectrum Analysis

Wireless networks operate in environments with electrical and radio frequency devices that can interfere with network communications. Microwave ovens, cordless phones, and even adjacent Wi-Fi networks are all potential sources of continuous or intermittent interference. The HTML-based spectrum analysis software modules on APs that support this feature examine the RF environment in which the Wi-Fi network is operating, identify interference and classify its sources. An analysis of the results quickly isolate issues with packet transmission, channel quality, and traffic congestion caused by contention with other devices operating in the same band or channel.

AP radios that gather spectrum data but do not service clients are called Spectrum Monitor (SM). Each SM scans and analyzes the spectrum band used by the radio (2.4 GHz or 5 GHz) of the SM. An AP radio in *hybrid AP* mode continues to serve clients as an access point while analyzing spectrum analysis data for the channel the radio uses to serve clients. You can record data for both types of spectrum analysis devices, save that data, and then play it back for later analysis.

Topics in this chapter include:

- Understanding Spectrum Analysis on page 900
- Creating Spectrum Monitors and Hybrid APs on page 905
- Spectrum Analysis Tasks on page 907
- Configuring Spectrum Analysis Dashboards on page 911
- Customizing Spectrum Analysis Graphs on page 913
- Working with Non-Wi-Fi Interferers on page 927
- Understanding Spectrum Analysis Session Log on page 928
- Viewing Spectrum Analysis Data on page 928
- Recording Spectrum Analysis Data on page 929

Understanding Spectrum Analysis

Single-radio mesh APs do not support the spectrum analysis feature; if an AP radio has a virtual AP carrying mesh backhaul traffic, no other virtual AP on that radio can be configured as a spectrum monitor. However, dual-radio mesh APs can have the client access radio configured as a Spectrum monitor or hybrid AP while the other radio supports mesh backhaul traffic.

This section describes the following topics:

- Device Support for Spectrum Analysis on page 900
- Viewing Spectrum Analysis
- Spectrum Analysis Clients
- Hybrid AP Channel Changes
- Hybrid APs Using Mode-Aware ARM

Device Support for Spectrum Analysis

The table below lists the AP models that support the spectrum analysis feature.

Table 161: *Device Support for Spectrum Analysis*

Device	Configurable as a Spectrum Monitor	Configurable as a Hybrid AP
570 Series	Yes	Yes
550 Series	Yes	Yes
530 Series	Yes	Yes
510 Series	Yes	Yes
500 Series	Yes	Yes
AP-518	Yes	Yes
AP-505H	Yes	Yes
320 Series	Yes	Yes
270 Series	Yes	Yes
220 Series	Yes	Yes
210 Series	Yes	Yes
AP-207	No	No
AP-205 / AP-205H	Yes	Yes
AP-204	Yes	Yes
AP-203H / AP-203R / AP-203RP	No	No
AP-115	Yes	Yes
AP-114	Yes	Yes
AP-104	Yes	Yes
RAP-3WN Series	Yes	No

Viewing Spectrum Analysis

The radios on groups of APs can be converted to dedicated spectrum monitors or hybrid APs via the 802.11A and 802.11G radio profiles of the AP group. Individual APs can also be converted to spectrum monitors through the spectrum override profile of the AP.



The spectrum analysis feature requires the RF Protect license. To convert an AP to a spectrum monitor or hybrid AP, you must have an AP license *and* an RFProtect license for each AP on that managed device.

The **Spectrum Analysis** tab of the **Diagnostics** > **Tools** in the WebUI includes the **Spectrum Dashboards** , **Spectrum Monitors**, and **Session Log** windows.

- **Spectrum Monitors**: this window displays a list of active spectrum monitors and hybrid APs streaming data to your client, the radio band the device is monitoring, and the date and time the SM or hybrid AP was connected to your client. This window allows you to select the spectrum monitors or hybrid APs for which you want to view information, and release the connection between your client and any device you no longer want to view.
- **Session Log**: this tab displays activity for spectrum monitors and hybrid APs during the current browser session, including timestamps showing when the devices were connected to and disconnected from the client, and any changes to a hybrid APs monitored channel.
- **Spectrum Dashboards**: this window shows different user-customizable data charts for 2.4 GHz and 5 GHz spectrum monitor or hybrid AP radios. Table 162 below gives a basic description of each of the spectrum analysis graphs that can appear on the spectrum dashboard.



For more detailed information on these graphs, refer to Customizing Spectrum Analysis Graphs on page 913.

Table 162: *Spectrum Analysis Graphs*

Graph Title	Description	Update Interval
Active Devices	A pie chart showing the percentages and total numbers of each device type for all active devices. This graph has no set update interval; the graph automatically updates when values change. For details, see Active Devices on page 914.	-
Active Devices Trend	A line chart showing the numbers of up to five different types of Wi-Fi and non-Wi-Fi devices seen on selected channels during a specified time interval. This chart can show devices on multiple channels for a spectrum monitor, or the single monitored channel for a hybrid AP. For details, see Active Devices Trend on page 915.	5 seconds
Channel Metrics	This stacked bar chart shows the current relative quality, availability or utilization of selected channels in the 2.4 GHz or 5 GHz radio bands. This chart can show multiple channels for a spectrum monitor, or the single monitored channel for a hybrid AP. For details, see Channel Metrics on page 916.	5 seconds
Channel Metrics Trend	A line chart showing the relative quality or availability of selected channels in the 2.4 GHz or 5 GHz radio bands over a specified time interval. Spectrum monitors can show channel data for multiple channels, while a hybrid AP shows information only for its one monitored channel. For details, see Channel Metrics Trend on page 917.	5 seconds
Channel Utilization Trend	A line chart that shows the channel utilization for one or more radio channels, as measured over a defined time interval. Spectrum monitors can show data for multiple channels, while a hybrid AP shows utilization levels for its one monitored channel only. For details, see Channel Utilization Trend on page 918.	5 seconds

Graph Title	Description	Update Interval
Device Duty Cycle	A stacked bar chart showing the percent of each channel in the radio frequency band of spectrum monitor used by a Wi-Fi AP or any other device type detected by the spectrum monitor. The Device Duty Cycle chart for a hybrid AP only shows data for the one channel monitored by the hybrid AP. For details, see Device Duty Cycle on page 919.	5 seconds
Devices vs Channel	A stacked bar chart showing the total numbers of each device type detected on each channel in the radio frequency band of the spectrum monitor. The Devices vs Channel chart for a hybrid AP only shows data for the one channel monitored by the hybrid AP. For details, see Devices vs Channel on page 920.	5 seconds
FFT Duty Cycle	Fast Fourier Transform, or **FFT**, is an algorithm for computing the frequency spectrum of a time-varying input signal. This line chart shows the FFT duty cycle, which represents the percent of time a signal is broadcast on the specified channel or frequency. Spectrum monitors can show data for multiple channels, while a hybrid AP shows information only for its one monitored channel. For details, see FFT Duty Cycle on page 921.	1 second
Interference Power	This chart shows information about Wi-Fi interference, including the Wi-Fi noise floor, and the amount of adjacent channel interference from cordless phones, bluetooth devices and microwaves. Spectrum monitors can show interference power data for multiple channels, while a hybrid AP shows information only for its one monitored channel. For details, see Interference Power on page 922.	5 seconds
Quality Spectrogram	This plot shows quality statistics for selected range of channels or frequencies as determined by the current noise floor, non-Wi-Fi (interferer) utilization and duty-cycles and certain types of retries. This chart can also be configured to show channel availability, the percentage of each channel that is unused and available for additional traffic. Spectrum monitors can show data for multiple channels, while a hybrid AP shows information only for its one monitored channel. For details, see Quality Spectrogram on page 923.	5 seconds
Real-Time FFT	Fast Fourier Transform, or **FFT**, is an algorithm for computing the frequency spectrum of a time-varying input signal. This line chart shows the power level of a signal on the channels or frequencies monitored by a spectrum monitor radio. Spectrum monitors can show data for multiple channels, while a hybrid AP shows information only for its one monitored channel. For details, see Real-Time FFT on page 924.	1 second
Swept Spectrogram	This plot displays FFT power levels For details, see or the FFT duty cycle for a selected channel or frequency, as measured during each time tick. Spectrum monitors can show data for multiple channels, while a hybrid AP shows information only for its one monitored channel. For details, see Swept Spectrogram on page 925.	1 second

Spectrum Analysis Clients

The maximum number of spectrum monitor radios and hybrid AP radios on a stand-alone controller is limited only by the number of APs on that stand-alone controller. If desired, you can configure every

radio on an AP that supports the Spectrum Analysis feature as a spectrum device. A dual-radio AP can operate as two spectrum devices, because each radio can be individually configured as a spectrum monitor or hybrid AP.

A spectrum analysis client can simultaneously access data from up to four individual spectrum device radios. Each spectrum device radio, however, can only be connected to a single client WebUI.

When you select a specific spectrum monitor or hybrid AP radio to stream data to your client, the stand-alone controller first verifies the device is not subscribed to some other client. Once the SM or hybrid AP radio has been verified as available, the SM or hybrid AP establishes a connection to the client and begins sending spectrum analysis data either every second or every five seconds, depending on the type of data being requested. Each client may select up to twelve different spectrum analysis charts and graphs to appear in the spectrum dashboard.

A stand-alone controller can support up to 22 active WebUI connections. If spectrum analysis clients are simultaneously viewing data for more than 22 WebUI connections, any additional WebUI requests are refused until some clients close their WebUI browser sessions.

When you finish reviewing data from an SM or hybrid AP, you should disconnect the device from your spectrum client. Do not forget this important step—no other user can access data from that spectrum monitor or hybrid AP until you release your subscription. Note, however, that when you disconnect a spectrum monitor from your client, *the AP continues to operate as a spectrum monitor* until you return it to AP mode by removing the local spectrum override, or by changing the mode parameter in the 802.11A or 802.11G radio profile from spectrum-mode back to AP-mode.



A spectrum monitor or hybrid AP automatically disconnects from a client when you close the browser window you used to connect the spectrum monitor to your client. However, if you use Internet Explorer and have multiple instances of an Internet Explorer browser open, the data-streaming connection to the spectrum monitor or hybrid AP is not released until 60 seconds after you close the spectrum client browser window. During this 60-second period, the spectrum monitor is still connected to the client.

When a spectrum monitor or hybrid AP is not subscribed to any client, it still performs all classification tasks and collect all necessary channel lists and device information. You can view classification, device, and channel information for any active spectrum monitor or hybrid AP via the command-line interface of the stand-alone controller, regardless of whether or not that device is sending real-time spectrum data to another client WebUI.

Individual spectrum analysis graphs and charts are explained in detail in Customizing Spectrum Analysis Graphs on page 913.

Hybrid AP Channel Changes

By default, a hybrid AP only monitors the channel specified in its 802.11A or 802.11G radio profile for spectrum interference. If you want to change the channel monitored by a hybrid AP, you must edit the channel setting in those profiles. However, there are other ArubaOS features that may automatically change the channels on hybrid APs. APs using DFS perform off-channel scanning to detect the presence of satellite and radar transmissions, and switch to a different channel if it detects that satellite or radar transmissions are present. APs using the ARM feature constantly monitor the network and automatically select the best channel and transmission power settings for that AP. If you manually change a channel monitored by a hybrid AP, best practices are to temporarily disable ARM, as ARM may automatically return the channel to its previous setting.

If a hybrid AP is using ARM or DFS, that hybrid AP may automatically move to a different channel in response to changes in the network environment. If a hybrid AP changes channels while it is connected to a spectrum analysis client, the hybrid AP updates the graphs in the spectrum dashboard to start

displaying spectrum data for the new channel, and sends a log message to the session log. For details on changing the channel monitored by a hybrid AP, see 2.4 GHz and 5 GHz Radio RF Management on page 676.

Hybrid APs Using Mode-Aware ARM

If a radio is configured as a hybrid AP and that AP is enabled with mode-aware ARM, the hybrid AP can change to an AM if too many APs are detected in the area. If ARM changes a hybrid AP to an Air Monitor, that AM does not provide spectrum data after the mode change. The AM unsubscribes from any connected spectrum analysis client, and sends a log message warning about the change. If mode-aware ARM changes the AM back to an AP, the hybrid AP does not automatically resubscribe back to the spectrum analysis client. The hybrid AP must manually resubscribed before it can appear in the **spectrum monitors** page of the client.

Creating Spectrum Monitors and Hybrid APs

Each stand-alone controller can support up to 22 active WebUI connections to spectrum monitor or hybrid AP radios. If you plan on using spectrum monitors or hybrid APs as a permanent overlay to constantly monitor your network, you should create a separate AP group for these devices. If you plan on temporarily converting campus APs to spectrum monitors, best practices are to use the spectrum local override profile to convert an AP to a spectrum monitor.

This section describes the following tasks for converting regular APs into hybrid APs or spectrum monitors.

- Converting APs to Hybrid APs on page 905
- Converting AP to Spectrum Monitor on page 906
- Converting Group of APs to Spectrum Monitors on page 907

Converting APs to Hybrid APs

You can convert a group of regular APs into a group of hybrid APs by selecting the **spectrum monitor** option in the 802.11A and 802.11G radio profiles of the AP group. Once you have enabled the spectrum monitor option, all APs in the group that support the spectrum monitoring feature start to function as hybrid APs. If any AP in the group does *not* support the spectrum monitoring feature, that AP continues to function as a standard AP, rather than a hybrid AP.



The spectrum monitoring option in the 802.11A and 802.11G radio profiles only affects APs in ap-mode. Devices in am-mode (Air Monitors) or sm-mode (Spectrum Monitors) are not affected by enabling this option.

If you want to convert a individual AP (and not an entire AP group) to a hybrid AP, you must create a new 802.11A or 802.11G radio profile, enable the **spectrum monitor** option, then reassign that AP to the new profile. For additional information see Creating and Editing Mesh High-Throughput SSID Profiles on page 757 for details on how to create a new 802.11A or 802.11G radio profile, then assign an individual AP to that profile.



If the spectrum local-override profile on the stand-alone controller that terminates the AP contains an entry for a hybrid AP radio, that entry overrides the mode selection in the 802.11A or 802.11G radio profile, and the AP operates as a spectrum monitor, not as a hybrid AP. You must remove any spectrum local override for an AP to allow the device to operate as a hybrid AP. For further details on editing a spectrum local override, see Converting AP to Spectrum Monitor on page 906.

The following procedure converts an AP group into hybrid APs:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration** > **AP Groups**.
2. Select an **AP Group**.
3. Click **Radio** tab for the selected AP group.
4. Under **Basic** > **Radio mode**, select **spectrum-mode** under either 2.4 GHz or 5 GHz.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

To convert a group of APs via the command-line interface, access the CLI in config mode and issue the following commands, where profile is the name of the 802.11A or 802.11G radio profile used by the group of APs you want to convert to hybrid APs:

```
rf dot11a-radio-profile <profile> spectrum-monitoring
rf dot11g-radio-profile <profile> spectrum-monitoring
```

Converting AP to Spectrum Monitor

There are two ways to change a radio on an individual AP or AM into a spectrum monitor. You can assign that AP to a different 802.11A and 802.11G radio profile that is already set to spectrum mode, or you can temporarily change the AP into a spectrum monitor using a local spectrum override profile. When you use a local spectrum override profile to override the mode setting of an AP, that AP begins to operate as a spectrum monitor, but remains associated with its previous 802.11A and 802.11G radio profiles. If you change any parameter (other than the overridden **mode** parameter) in the 802.11A or 802.11G radio profiles of the spectrum monitor, the spectrum monitor immediately updates with the change. When you remove the local spectrum override, the spectrum monitor reverts back to its previous mode, and remains assigned to the same 802.11A and 802.11G radio profiles as before.

The spectrum local override profile overrides the **mode** parameter in the 802.11A or 802.11G radio profile, changing it from ap-mode or am-mode to spectrum-mode, while allowing the spectrum monitor to continue to inherit all other settings from its 802.11A or 802.11G radio profiles. When the spectrum local override is removed, the AP automatically reverts to its previous mode as defined it its 802.11A or 802.11G radio profile settings. If you use the local override profile to change an AP radio to a spectrum monitor, you must do so by accessing the WebUI or CLI of the stand-alone controller that terminates the AP.

The following procedure converts an individual AP using the spectrum local override profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration** > **System**.
2. Select **Profiles** tab.
3. Select **Spectrum Local Override**.
4. In the **Spectrum Local Override Profile**, click **+**.
5. In the **ap_name**, enter the name of an AP whose radio you want to configure as a spectrum monitor. The AP names are case-sensitive. Any extra spaces before or after the AP name prevents the AP from being correctly added to the override list.

6. If your AP has multiple radios or a single dual-band radio, click the **spectrum_band** drop-down list and select the spectrum band you want that radio to monitor: **2.4 GHz** or **5 GHz**.
7. Click **OK** to add that radio to the **Override Entry** list.
8. Repeat steps 4 through 7 to convert other AP radios to spectrum monitors, if required.
9. To remove spectrum monitor from the override entry list, select that radio name in the override entry list, then click **Delete**.
10. Click **Submit**.
11. Click **Pending Changes**.
12. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

Access the CLI in config mode and issue the following command to convert an individual AP using the spectrum local override profile:

```
ap spectrum local-override override ap-name <ap-name> spectrum-band 2ghz|5ghz
```

Converting Group of APs to Spectrum Monitors

When you convert a group of APs to spectrum monitors using their 802.11A or 802.11G radio profiles, all AP radios associated with that profile stop serving clients and act as spectrum monitors only. Therefore, before you convert an entire group of APs to spectrum monitors, be sure that none of the APs are currently serving clients, as that may temporarily interrupt service to those clients.



If you use an 802.11A or 802.11G radio profile to create a group of spectrum monitors, all APs in any AP group referencing that radio profile are set to spectrum mode. Therefore, best practices are to create a new 802.11A or 802.11G radio profile just for spectrum monitors, using the following CLI commands:

```
ap-name <ap name> dot11a-radio-profile <profile-name>ap-name <ap name>
dot11g-radio-profile <profile-name>
```

If you want to set an existing 802.11A or 802.11G radio profile to spectrum mode, verify that no other AP group references that radio profile, using the following CLI commands:

```
show references rf dot11a-radio-profile <profile-name>show references rf
dot11g-radio-profile <profile-name>
```

The following procedures convert an AP group into Spectrum mode:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration** > **AP Groups**.
2. Select an **AP Group**.
3. Click **Radio** tab for the selected AP group.
4. Under **Basic** > **Radio mode**, select **spectrum-mode** under either 2.4 GHz or 5 GHz.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

Access the CLI in config mode and issue the following commands, where **<profile>** is the 802.11A or 802.11G radio profile used by the AP group.

```
rf dot11a-radio-profile <profile> mode spectrum-mode
rf dot11g-radio-profile <profile> mode spectrum-mode
```

Spectrum Analysis Tasks

A spectrum analysis client is any laptop or desktop computer that can access a stand-alone controller WebUI and receive streaming data from individual spectrum monitors or hybrid APs. Once you have configured one or more APs to operate as a spectrum monitor or hybrid AP, navigate to the **Managed Network** node hierarchy from the Mobility Master WebUI, and use the **Spectrum Monitors** window to identify the spectrum devices you want to actively connect to the spectrum analysis client.



The Spectrum Analysis option is not available if the license is not enabled or present.

The following sections explain the various tasks related to spectrum analysis as per the needs of your individual network:

- Obtaining Managed Device Node Details on page 908
- Connecting Spectrum Devices to Spectrum Analysis Client on page 908
- Viewing Connected Spectrum Analysis Devices on page 909
- Disconnecting Spectrum Device on page 909
- Verifying Spectrum Monitors Support for One Client per Radio on page 910
- Converting a Spectrum Monitor Back to an AP or Air Monitor on page 910
- Troubleshooting Browser Issues on page 910
- Loading a Spectrum View on page 911
- Understanding Spectrum Analysis Syslog Messages on page 911
- Playing a Recording in the RFPlayback Tool on page 911

Obtaining Managed Device Node Details

Before you log in to the Spectrum Analysis window, you must obtain the managed device node where the AP or spectrum monitor is terminated by using one of the following methods:

- In the **Managed Network** node hierarchy, navigate to **Configuration** > **Access Points** > **Campus APs** tab to obtain the managed device node details from the **Switch IP** column.
- In the **Managed Network** node hierarchy, navigate to **Dashboard** > **Infrastructure** > **Access Devices** tab to obtain the managed device node details from the **Active Controller** column.

Connecting Spectrum Devices to Spectrum Analysis Client

To connect one or more spectrum devices to your client, log in to the managed device obtained from the **Managed Network** node hierarchy (as described in Obtaining Managed Device Node Details on page 908):

1. In the **Managed Device** node hierarchy, navigate to **Diagnostics** > **Tools** > **Spectrum Analysis**.
2. Click **Launch** to launch the Spectrum Analysis tool.
3. Click the **Spectrum Monitors** tab in the new window.
4. Click **Add**. A table appears, displaying a list of spectrum analysis devices, sorted by name. Single-radio spectrum devices have a single entry in this table, and dual-radio spectrum devices have two entries: one for each radio. This table displays the following data for each radio.

Table 163: *Spectrum Device Selection Information*

Table Column	Description
AP	Name of the AP whose radio you want to convert to a spectrum monitor. AP names are case sensitive. This column includes the following icons:  Radio is operating as a spectrum monitor.  Radio is operating as a hybrid AP with spectrum enabled.
Band	The frequency band currently used by the radio. This value can be either **2.4 GHz** or **5 GHz**.
Model	AP model type.
AP Group	Name of the AP group to which the spectrum monitor is currently associated.
Mode	This column indicates the type of spectrum analysis device: ▪ **Spectrum Monitor**: AP is in spectrum monitor mode. ▪ **Access Point**: AP is configured as an access point but with spectrum monitoring enabled Hybrid AP.
Availability for Connection	Indicates if the AP is available to send spectrum analysis data to the client. Possible options ares: ▪ **Available, 2.4 GHz**: the radio is available to send spectrum analysis data on the 2.4 GHz frequency band. ▪ **Available, 5 GHz**: the radio is available to send spectrum analysis data on the 5 GHz frequency band. ▪ **Available, Dual Band**: the radio is available and is capable of sending spectrum analysis data on either the 2.4 GHz or the 5 GHz frequency bands. ▪ **Available, current channel - <channel>**: the AP radio is in hybrid mode and can display spectrum analysis data for the single specified channel only. ▪ **Not available:** an AP may not be available because it is currently sending spectrum analysis data to another client.

5. Click the table entry for a spectrum monitor radio, then click **Connect**.
6. Repeat steps 3-4 to connect additional devices, if desired.

Viewing Connected Spectrum Analysis Devices

Once you have connected one or more spectrum monitors or hybrid APs to your Spectrum Analysis client, the **Diagnostics** > **Tools** > **Spectrum Analysis** > **Spectrum Monitors** window displays a table of currently connected spectrum devices. This table includes the name of each spectrum monitor or hybrid AP and its current radio band (2.4 GHz or 5 GHz):

To view a list of connected spectrum devices via the command-line interface, issue the **show ap spectrum monitors** command.

Disconnecting Spectrum Device

A spectrum monitor or hybrid AP can send spectrum analysis data to only one client at a time. When you are done viewing data for a spectrum device, you should release the subscription of your client to

that spectrum device and allow other clients to view data from that device. A spectrum monitor or hybrid AP automatically disconnects from your client when you close the browser window used to connect the spectrum device your client.

To manually disconnect a spectrum monitor or hybrid APs, log in to the managed device obtained from the **Managed Network** node hierarchy (as described in Obtaining Managed Device Node Details on page 908):

1. In the **Managed Device** node hierarchy, navigate to **Diagnostics** > **Tools** > **Spectrum Analysis**.
2. Click **Launch** to launch the Spectrum Analysis tool.
3. Click the **Spectrum Monitors** tab in the new window.
4. Each table entry in the **Currently Connected** table includes a **Disconnect** link to release the connection of the client to that spectrum monitor. Identify the table entry for the spectrum monitor you want to release then click **Disconnect**.
5. A popup window asks you to confirm that you want to disconnect the spectrum monitor from the spectrum analysis client. Click **OK**. The spectrum monitor disconnects from the client and the device's entry is removed from the **Currently Connected** table.

When you disconnect a spectrum device from your client, the AP continues to operate as a spectrum monitor or hybrid AP until you return the device to AP mode by removing the local spectrum override, or by changing the mode parameter in the 802.11A or 802.11G radio profile of the AP from spectrum-mode to AP-mode.



If you are use Internet Explorer with multiple instances of the Internet Explorer browser open, and you close the spectrum browser window without manually disconnecting the spectrum device, the stand-alone controller does not release the data streaming connection to a spectrum monitor until 60 seconds after you close the spectrum client browser window. During this 60-second period, the spectrum monitor is still connected to the client.

Verifying Spectrum Monitors Support for One Client per Radio

Each spectrum monitor radio can only send information to one client at a time. If you log into a stand-alone controller and the spectrum monitor dashboard does not display any data for the selected radio, another user may be logged in to the stand-alone controller at that time. Note that dual-radio spectrum monitors may be accessed by two clients; one client for each radio.

Converting a Spectrum Monitor Back to an AP or Air Monitor

If want to convert a spectrum monitor radio back to AP or AM mode but the radio still comes up as a spectrum monitor, access the command-line interface and see if that spectrum monitor appears in the output of the **show ap spectrum local-override** command. If the spectrum monitor does appear in the local override profile table, issue the command **ap spectrum local-override no override ap-name <apname> spectrum-band <spectrum-band>** to remove the local override for that spectrum monitor and return the radio to AP or AM mode.

Troubleshooting Browser Issues

If you access the spectrum analysis dashboard using the Safari 5.0 browser, clicking the backspace button may return you to the previous browser screen. Avoid using the backspace button when changing dashboard view names or chart options.

If you are recording spectrum analysis data or playing back a spectrum analysis recording using a Mac client, do not minimize the browser window while the recording is in progress.

Loading a Spectrum View

Saved spectrum view preferences may not be backwards compatible with the spectrum analysis dashboard in earlier versions of ArubaOS. If you downgrade to an earlier version of ArubaOS and your client is unable to load a saved spectrum view in the spectrum dashboard, access the CLI and issue the command **ap spectrum clear-webui-view-settings** to delete the saved spectrum views and display default view settings in the spectrum dashboard.

Understanding Spectrum Analysis Syslog Messages

The spectrum analysis feature can send four different types of syslog messages: wifi add, wifi delete, non-wifi add, and non-wifi delete. All messages are in the wireless category at the syslog severity level NOTICE.

The four syslog message types appear in the following formats:

- AM: Spectrum: new wifi device found = [addr:%s] SSID = [ssid:%s] BSSID [bssid_str:%s] DEVICE ID [did:%d]
- AM: Spectrum: deleting wifi device = [addr:%s] SSID = [ssid:%s] BSSID [bssid_str:%s] DEVICE ID [did:%d]
- AM: Spectrum: new non-wifi device found = DEVICE ID [did:%u] Type [dytpe:%s] Signal [sig:%u] Freq [freq:%u]KHz Bandwidth [bw:%u]KHz
- AM: Spectrum: deleting non-wifi device = DEVICE ID [did:%d] Type [dtype:%s]

Playing a Recording in the RFPlayback Tool

The Aruba RFPlayback tool is periodically updated to support improvements to the ArubaOS Spectrum Analysis feature. The RFPlayback tool can play spectrum recordings created in the same version of ArubaOS or earlier releases. If the RFPlayback tool cannot load a newer recording, you may need to download a more recent version of the tool from the Aruba website.

Configuring Spectrum Analysis Dashboards

Once you have connected spectrum monitors to your spectrum analysis client, you can begin to monitor spectrum data in the spectrum analysis dashboards. There are three predefined sets of dashboard views, **View 1**, **View 2** and **View 3**. View 1 displays the Real-Time FFT, FFT Duty-Cycle and Swept Spectrogram graphs by default, and Views 2 and 3 display the Swept Spectrogram and Quality Spectrogram charts, and the Channel Summary and Active Devices tables.



Spectrum Analysis dashboards are available only on stand-alone controllers.

Each chart in the dashboard can be replaced with other chart types, or reconfigured to show data for a different spectrum monitor. Once you have configured a dashboard view with different settings, you can rename that dashboard view to better reflect its new content.

The following sections explain how to customize your Spectrum Analysis dashboard to best suit the needs of your individual network:

- Selecting Spectrum Monitor on page 912
- Changing Graphs within Spectrum View on page 912
- Renaming Spectrum Analysis Dashboard View on page 912
- Saving Dashboard View on page 913
- Resizing an Individual Graph on page 913

Selecting Spectrum Monitor

When you first log in to the **Spectrum Analysis** dashboard from the Managed Device WebUI, it displays blank charts. You must identify the spectrum monitor whose information you want to view before the graphs display any data.

To identify the spectrum monitor radio whose data you want to display in the **Spectrum Analysis** dashboard, log in to the managed device obtained from the **Managed Network** node hierarchy (as described in **Obtaining Managed Device Node Details** section of Spectrum Analysis Tasks on page 907):

1. In the **Managed Device** node hierarchy, navigate to **Diagnostics** > **Tools** > **Spectrum Analysis**.
2. Click **Launch** to launch the Spectrum Analysis tool.
3. Click the **Spectrum Monitors** tab in the new window.
4. Click **Add**.
5. Select a spectrum monitor from the list and click **Connect**.

After you have selected the initial spectrum monitor or hybrid AP for a graph, you can display data for a different spectrum device at any time by clicking the down arrow by the device name in the chart titlebar and selecting a different connected spectrum monitor or hybrid AP.

Changing Graphs within Spectrum View

To replace an existing graph with any other type of graph or chart, log in to the managed device obtained from the **Managed Network** node hierarchy (as described in **Obtaining Managed Device Node Details** section of Spectrum Analysis Tasks on page 907):

1. In the **Managed Device** node hierarchy, navigate to **Diagnostics** > **Tools** > **Spectrum Analysis**.
2. Click **Launch** to launch the Spectrum Analysis tool.
3. Click the **Spectrum Dashboards** tab.
4. From **Spectrum Dashboards** window, click one of the view names at the top of the window to select the dashboard layout with the graph you want to change.
5. Click the down arrow at the far right end of the graph title bar to display a drop-down list of chart options.
6. Click **Replace With** to display a list of available graphs.
7. Click the name of the new graph you want to display.

Renaming Spectrum Analysis Dashboard View

You can rename any of the three spectrum analysis dashboard views at any time. However, renaming a view does not save its settings. (For details on saving a spectrum dashboard view, refer to Saving Dashboard View on page 913.)

To rename a Spectrum Analysis Dashboard view, log in to the managed device obtained from the **Managed Network** node hierarchy (as described in **Obtaining Managed Device Node Details** section of Spectrum Analysis Tasks on page 907):

1. In the **Managed Device** node hierarchy, navigate to **Diagnostics** > **Tools** > **Spectrum Analysis**.
2. Click **Launch** to launch the Spectrum Analysis tool.
3. Click the **Spectrum Dashboards** tab.
4. Click the down arrow to the right of the dashboard view you want to rename.
5. Select **Rename**.

6. The **Dashboard Name** popup window appears. Enter a new name for the dashboard view, then click **OK**.

Saving Dashboard View

You can select different graphs to display in a dashboard view, but these changes are not saved unless you save that view. Dashboard views, (like the spectrum analysis profile and spectrum local-override profile) are all local configurations that must be configured on each stand-alone controller.

The following procedures save a dashboard view:

1. After selecting the graphs you want to appear in the view, click **Save Spectrum Views** at the top of the window.
2. The **Spectrum View saved successfully** confirmation window appears when the spectrum view has been saved.



If you change graphs in a spectrum view but do not save your settings, you are prompted to save or cancel your changes when you close the spectrum dashboard browser window.

Resizing an Individual Graph

The left side of the title bar for each graph includes a resizing button on that allows you to expand a graph for easier viewing. Click this button to expand the selected graph to the size of the full window and display the **Options** pane, which allows you to change the current display options for that graph. (Configuration options are described in Spectrum Analysis Graph Configuration Options on page 913). To close the options pane if you have not made any changes to the graph, click **Close** at the bottom of the **Options** pane *or* click the resize button again to return the graph to its original size. To save any changes to the graph, click **OK** to save your settings and close the **Options** pane.

Customizing Spectrum Analysis Graphs

Each Spectrum Analysis graph can be customized to display or hide selected data types. To view the available options for a graph type, log in to the managed device obtained from the **Managed Network** node hierarchy in the Mobility Master (as described in **Obtaining Managed Device Node Details** section of Spectrum Analysis Tasks on page 907):

1. In the **Managed Device** node hierarchy, navigate to **Diagnostics** > **Tools** > **Spectrum Analysis**.
2. Click **Launch** to launch the Spectrum Analysis tool.
3. Click the **Spectrum Dashboards** tab in the new window.
4. Click the down arrow at the end of the title bar for the graph you want to configure.
5. Select **Options**. The **Options** window appears to the right of the graph.
6. From the **Options** window, configure graph settings described in Spectrum Analysis Graph Configuration Options on page 913.
7. When you are done, click **Close** at the bottom of the **Options** window to hide the options window.
8. Click **Save Spectrum Views** at the top of the window to save your new settings.

Spectrum Analysis Graph Configuration Options

The following sections describe the customizable parameters and the default settings for each spectrum analysis graph.

This section contains the following topics:

- Active Devices on page 914
- Active Devices Trend on page 915
- Channel Metrics on page 916
- Channel Metrics Trend on page 917
- Channel Utilization Trend on page 918
- Device Duty Cycle on page 919
- Devices vs Channel on page 920
- FFT Duty Cycle on page 921
- Interference Power on page 922
- Quality Spectrogram on page 923
- Real-Time FFT on page 924
- Swept Spectrogram on page 925

Active Devices

This graph appears as a pie chart showing the percentages and total numbers of each device type for all active devices seen by the spectrum monitor or hybrid AP radio. This chart is useful for determining which types of devices are sending signals on the specified radio band or channel. The Active Devices graphs for spectrum monitors can be configured to show data for several different device types on a single radio channel or range of channels. Active Devices graphs for hybrid APs can show data for the single monitored channel only.

When you hover your mouse over any section of the pie chart, a tooltip displays the percentage and number of active devices classified into that device type.

Click the down arrow in the upper right corner of this chart then click the **Options** menu to access the configuration settings for the Active Devices graph. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 164: *Active Devices Graph Options*

Parameter	Description
Band	Radio band displayed in this graph (2.4 GHz or 5 GHz)
Channel numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11AC include an additional **80 MHz** option for very-high-throughput channels.
Channel Range	For graphs created by spectrum monitors, specify a channel range to determine which channels appears in this graph. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the graph. This graph displays all channels within the radio band of the spectrum monitor by default. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.

Active Devices Trend

The Active Devices Trend chart is a line chart that shows the numbers of Wi-Fi and non-Wi-Fi devices seen on each radio channel during the displayed time interval. When you hover your mouse over any line in the chart, a tooltip displays the number of active devices for the selected device type.

An Active Devices Trend chart created by a hybrid AP displays data for the single channel monitored by that device. For spectrum monitors, the Active Devices Trend chart can display values for up to five different channels and device types. These graphs show the following data by default:

- For SMs on the 2.4 GHz radio band, Wi-Fi APs on channel 1, 6, and 11.
- For SMs on the 5 GHz band, Wi-Fi APs on channel 36, 40, and 44.

Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access the Active Devices Trend configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 165: *Active Devices Trend Options*

Parameter	Description
Band	Radio band displayed in this graph (2.4 GHz or 5 GHz).
Show Trend for Last	Amount of elapsed time for which this chart should display data.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the Channel Numbering drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11AC include an additional **80 MHz** option for very-high-throughput channels.
Show lines for these channels	The Active Devices Trend chart can display values for up to five different device types on different channels for a spectrum monitor, or a single device type for a hybrid AP. To choose which type of data each line should represent, click the channel number drop-down list and select a channel within the radio band, then click the device type drop-down list and select one of the following device types. ■ WiFi (AP) ■ Microwave *(This option is only available for 2.4 GHz radios)* ■ Bluetooth *(This option is only available for 2.4 GHz radios)* ■ Fixed Freq (Others) ■ Fixed Freq (Cordless Phones) ■ Fixed Freq (Video) ■ Fixed Freq (Audio) ■ Freq Hopper (Others) ■ Freq Hopper (Cordless Network) ■ Freq Hopper (Cordless Base) ■ Freq Hopper Xbox *(This option is only available for 2.4 GHz radios)* ■ Microwave (Inverter) *(This option is only available for 2.4 GHz radios)* ■ Generic Interferer

Parameter	Description
	Select the check box beside each channel and device entry to show that information on the chart, or deselect the check box to hide that information. For more information on non-Wi-Fi device types detected by a spectrum monitor, see Working with Non-Wi-Fi Interferers on page 927.

Channel Metrics

This stacked bar chart can show one of three different types of channel metrics; **channel utilization**, **channel availability**, or **channel quality**.

This chart displays channel utilization data by default, showing both the percentage of each monitored channel that is currently being used by Wi-Fi devices, and the percentage of each channel being used by non-Wi-Fi devices and 802.11 ACI.



ACI refers to the interference on a channel created by a transmitter operating in an adjacent channel. A transmitter on a nonadjacent or partially overlapping channel may also cause interference, depending on the transmit power of the interfering transmitter and the distance between the devices. In general, ACI may be caused by a Wi-Fi transmitter or a non-Wi-Fi interferer. However, whenever the term ACI appears in Spectrum Analysis graphs, it refers to the ACI caused by Wi-Fi transmitters. The channel utilization option in the Channel Metrics Chart shows the percentage of the channel utilization due to both ACI and non-Wi-Fi interfering devices. Unlike the ACI shown in the **Interference Power** chart, the ACI shown in this graph indicates the percentage of channel time that is occupied by ACI or unavailable for Wi-Fi communication due to ACI.

The Channel Metrics graph can also show channel availability, the percentage of each channel that is available for use, or display the current relative quality of selected channels in the 2.4 GHz or 5 GHz radio bands. Spectrum monitors can display data for all channels in their selected band. Hybrid APs display data for their one monitored channel only.

In the spectrum analysis feature, channel quality is a relative measure that indicates the ability of the channel to support reliable Wi-Fi communication. Channel quality, which is represented as a percentage in this chart, is a weighted metric derived from key parameters that can affect the communication quality of a wireless channel, including noise, non-Wi-Fi (interferer) utilization and duty-cycles, and certain types of retries. Note that channel quality is not directly related to Wi-Fi channel utilization, as a higher quality channel may or may not be highly used.

When you hover your mouse over any bar in the chart, a tooltip displays the metric value for that individual channel. The example below shows that 61% of channel 3 is being consumed by non-Wi-Fi devices and 802.11 adjacent channel interference.

Click the down arrow in the upper right corner of this chart then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 166: *Channel Metrics Options*

Parameter	Description
Band	Radio band displayed in this graph. For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5 GHz Low, 5 GHz Center or 5 GHz High** to display data for that portion of the 5 GHz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.

Parameter	Description
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11AC include an additional **80 MHz** option for very-high-throughput channels.
Channel Range	For graphs created by spectrum monitors, specify a channel range to determine which channels appear in this graph. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the graph. This graph displays all channels within the radio band of the spectrum monitor by default. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.
Display Mode	▪ Select **Channel Quality** to show the relative quality of the channel. Channel Quality is a weighted metric derived from key parameters which include noise, non-Wi-Fi (interferer) utilization and duty-cycles, and certain types of retries. ▪ Select **Channel Availability** to show the percentage of the channel that is unused and available for additional Wi-Fi traffic. ▪ Select **Channel Utilization** to show both the percentage of the channel that is currently used by Wi-Fi devices, and the percentage of each channel that is being used by non-802.11 devices or 802.11 ACI.

Channel Metrics Trend

By default, this line chart shows the current relative quality of selected channels in the 2.4 GHz or 5 GHz radio bands over a period of time. The Channel Metrics Trend chart can also be configured to display trends for the current availability of selected channels, or the percentage of availability for those channels. Spectrum monitors can display data for up to five different channels. Hybrid APs display data for their one monitored channel only.



For more information on how the spectrum analysis feature determines the quality of a channel, see Channel Metrics on page 916.

When you hover your mouse over any line in the chart, a tooltip displays channel quality or availability data for that individual channel at the selected time.

Click the down arrow in the upper right corner of this chart then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboard.

Table 167: *Channel Metrics Trend Options*

Parameter	Description
Band	Radio band displayed in this graph (2.4 GHz or 5 GHz).

Parameter	Description
Show Trend for Last	The Channel Quality Trend chart shows channel quality or channel availability for the past 10 minutes by default. To view data for a different time range, click the **Show Trend for Last** drop-down list and select one of the following options: ■ 10 minutes ■ 30 minutes ■ 1 hour
Channel numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11AC include an additional **80 MHz** option for very-high-throughput channels.
Show Lines for These Channels	The Channel Quality Trend chart for a spectrum monitor can display channel quality, channel availability, or channel utilization values for up to five different channels on the selected radio band. Charts for hybrid APs can display data for the one channel monitored by that hybrid AP radio. To choose which type of data each line should represent on a chart for a spectrum monitor, click the **channel number** drop-down list and select a channel within the radio band, then click the second drop-down list and select either **Channel Quality**, or **Channel Availability**. Select the check box beside each channel entry to show that information on the chart, or deselect the check box to hide that information.

Channel Utilization Trend

The Channel Utilization Trend chart is a line chart that shows the percentage of total utilization on each channel over a time interval. The channel utilization includes the utilization due to Wi-Fi as well as utilization due to non-Wi-Fi interferers and ACI.



For additional information on how the spectrum analysis feature measures ACI, see Channel Metrics on page 916.

This graph can show data recorded for the last ten, thirty, or sixty minutes. Spectrum monitors display spectrum analysis data seen on all channels in the selected band, and hybrid APs display data from the one channel they are monitoring. When you hover your mouse over any line in the chart, a tooltip shows the percentage of the channel being utilized at the specified time.

Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 168: *Channel Utilization Trend Options*

Parameter	Description
Intervals	The Channel Utilization Trend chart shows channel quality or channel availability for the past 10 minutes by default. To view data for a different time range, click the **Intervals** drop-down list and select one of the following options: ■ 10 minutes ■ 30 minutes

Parameter	Description
	▪ 1 hour
Band	Radio band displayed in this graph (2.4 GHz or 5 GHz).
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11AC include an additional **80 MHz** option for very-high-throughput channels.
Show	To select individual channels you want to display on this chart, click the check box by a channel entry, then click the **channel** drop-down list to select the channel to display. To hide a channel, uncheck the check box by that channel number.

Device Duty Cycle

The Device Duty Cycle Chart is a stacked bar chart that shows the duty cycle of each device type on a channel. The duty cycle is the percentage of time each device type operates or transmits on that channel. Though Wi-Fi devices do not transmit if there is another Wi-Fi or non-Wi-Fi device active at that time, most non-Wi-Fi devices do not follow such a protocol for transmissions. Because these devices operate independently without regard to any other devices operating on the same channel, the total duty cycle of all device types may add up to more than 100% on a channel. For example, one or more video bridges may be active on a channel, each with a 100% duty cycle. The same channel may have a cordless transmitter with a 10% duty cycle and a microwave oven with a 50% duty cycle. In this example, the Device Duty Cycle chart shows all three device types with their respective duty cycle percentages.



A hybrid AP on a 20 MHz channel will see 40 MHz Wi-Fi data as non-Wi-Fi data.

Spectrum monitors display spectrum analysis data seen on all channels in the selected band, and hybrid APs display data from the one channel they are monitoring. The example below shows data from a spectrum monitor monitoring all channels in the 2.4 GHz band.

Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 169: *Device Duty Cycle Options*

Parameter	Description
Band	Radio band displayed in this graph. For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5 GHz Low**, **5 GHz Center** or **5 GHz High** to display data for that portion of the 5 GHz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.

Parameter	Description
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11AC include an additional **80 MHz** option for very-high-throughput channels.
Channel Range	For graphs created by spectrum monitors, specify a channel range to determine which channels appear in this graph. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the graph. This graph displays all channels within the radio band of the spectrum monitor by default. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.

Devices vs Channel

This stacked bar chart shows the current number of devices using each channel in the frequency band of the radio. This chart can show separate per-channel statistics for the numbers of Wi-Fi devices, cordless phones, bluetooth devices, microwaves, and other non-Wi-Fi devices.

If a device affects more than one channel, it is recorded as a device on all channels it affects. For example, if a 20 MHz Wi-Fi AP has a center frequency of 2437 MHz (channel 6) it is counted as a device on channels 3-9 because it affects all those channels. Similarly, if a channel-hopping device uses all channels within a frequency band, it is counted as a device on all channels in that band.

When you hover the mouse over any part of the chart, a tooltip shows the numbers of the device type currently using that channel.

Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 170: *Devices vs Channel Options*

Parameter	Description
Band	Radio band displayed in this graph. For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5 GHz Low, 5 GHz Center or 5 GHz High** to display data for that portion of the 5 GHz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11AC include an additional **80 MHz** option for very-high-throughput channels.
Channel Range	For graphs created by spectrum monitors, specify a channel range to determine which channels appear in this graph. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the graph.

Parameter	Description
	This graph displays all channels within the radio band of the spectrum monitor by default. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.

FFT Duty Cycle

The FFT Duty Cycle chart is a line chart that shows the duty cycle for each frequency bin. The width of the each frequency bin depends on the resolution bandwidth of the spectrum monitor. The spectrum analysis feature considers a frequency bin to be used if the detected power in that bin is at least 20 dB higher than the nominal noise floor on that channel. The FFT Duty Cycle provides a more granular view of the duty cycle per bin as opposed to the aggregated channel utilization reported in the Channel Metrics chart.



A hybrid AP on a 20 MHz channel will see 40 MHz Wi-Fi data as non-Wi-Fi data.

This chart can show the duty cycle over the last second, the maximum FFT duty cycle measured for all samples taken over the last N sweeps, and the greatest FFT duty cycle recorded since the chart was last reset.

Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 171: *FFT Duty Cycle Options*

Parameter	Description
Band	Radio band displayed in this graph. For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5 GHz Low**, **5 GHz Center** or **5 GHz High** to display data for that portion of the 5 GHz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11AC include an additional **80 MHz** option for very-high-throughput channels.
X-Axis	Select either **Channel** or **Frequency** to show the duty cycle for a range of channels or frequencies.
Channel Range	If you selected **Channel** in the **X-Axis** parameter, you must also specify a channel range to determine which channels appear in the x-axis of this chart. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the chart. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.

Parameter	Description
Y-Axis	Select either **Frequency** or **Channel** to show the duty cycle for a range of frequencies or channels.
Show	Select a check box to display that information on the FFT Duty Cycle chart. ■ **Duty Cycle:** The percentage of duty cycle the channel or frequency was actively used. ■ **Max Hold:** The maximum recorded percentage of active duty cycles for the channel frequency since the chart was last reset. To clear this setting, click the down arrow at the end of the title bar for this graph and select **Reset MaxHold**. ■ **Max of last sweeps:** This chart shows the maximum percentage of active duty cycles for the channel of frequency recorded during the last 10 sweeps, by default. To change the number of sweeps used to determine this value, enter a number from 2 to 20, inclusive. To clear this setting, click the down arrow at the end of the title bar for this graph and select **Reset MaxNSweep**.

Interference Power

The Interference Power chart displays various power levels of interest, including the Wi-Fi AP with maximum signal strength, noise, and interferer types with maximum signal strength. The ACI displayed in the Interference Power Chart is the ACI power level based on the signal strength(s) of the Wi-Fi APs on adjacent channels. A higher ACI value in Interference Power Chart does not necessarily mean higher interference, because the AP that is contributing to the maximum ACI may or may not be very actively transmitting data to other clients at all times. The ACI power levels are derived from the signal strength of the beacons.

This chart displays the noise floor of each selected channel in dBm. The noise floor of a channel depends on the noise figure of the RF components used in the radio, temperature, presence of certain types of interferers or noise, and the width of the channel. For example, in a clean RF environment, a 20 MHz channel has a noise floor around -95 dBm and a 40 MHz channel has a noise floor around -92 dBm. Certain types of fixed-frequency continuous transmitters such as video bridges, fixed-frequency phones, and wireless cameras typically elevate the noise floor seen by the spectrum monitor. Other interferers such as frequency-hopping phones, Bluetooth, and Xbox may not affect the noise floor of the radio. A Wi-Fi radio can only reliably decode Wi-Fi signals that are a certain dB above the noise floor. Therefore estimating and understanding the actual noise floor of the radio is critical to understanding the reliability of the RF environment.

The chart also includes information about the AP on each channel with the highest power level. You can hover your mouse over an AP on the chart to view the name, SSID, and current power level of the AP. The example below shows that the AP with the maximum power on channel 157 has the SSID **qa-ss**, and a power level of -55dBm.

Click the down arrow in the upper right corner of this chart then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 172: *Interference Power Options*

Parameter	Description
Band	Radio band displayed in this graph.

Parameter	Description
	For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5 GHz Low, 5 GHz Center or 5 GHz High** to display data for that portion of the 5 GHz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80MHz** option for very-high-throughput channels.
Channel Range	For graphs created by spectrum monitors, specify a channel range to determine which channels appear in this graph. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the graph. This graph displays all channels within the radio band of the spectrum monitor by default. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.

Quality Spectrogram

This plot shows the channel quality statistics for selected range of channels or frequencies. This chart can also be configured to show channel availability, the percentage of each channel that is unused and available for additional traffic.

Channel Quality is a weighted metric derived from key parameters which include noise, non-Wi-Fi (interferer) utilization and duty-cycles and certain types of retries. Quality levels are indicated by a range of colors between dark blue, which represents a higher channel quality, and red, which represents a lower channel quality. Channel availability is indicated by a range of colors between dark blue, which represents 100% channel availability, and red, which represents 0% availability.



For additional information on interpreting anAruba Spectrogram plot, see Swept Spectrogram on page 925.

The Spectrum Analysis Quality Spectrogram chart measures channel data each second, so after every 5-second sweep, the newest data appears as a thin colored line on the bottom of the chart. Older data is pushed up higher on the chart until it reaches the top of the spectrogram and ages out. The example below shows the Aruba Quality Spectrogram chart after it has recorded over 1,500 seconds of FFT data.

When you hover your mouse over any part of the spectrogram, a tooltip shows the devices the spectrum monitor detected on that frequency, the BSSID of the device (if applicable), the power level of the device in dBm, the time the device was last seen by the spectrum monitor, and the channels affected by the device.

The following table describes the other optional parameters you can use to customize the Quality Spectrogram. Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 173: *Quality Spectrogram Options*

Parameter	Description
Band	Radio band displayed in this graph. For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5GHz Low**, **5GHz Center** or **5GHz High** to display data for that portion of the 5GHz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80MHz** option for very-high-throughput channels.
X-Axis	Select either **Channel** or **Frequency** to show the quality spectrogram for a range of channels or frequencies.
Channel Range	Specify a channel range to determine which channels appear in the x-axis of this chart. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the chart. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.
Color-Map Range	*If this chart is configured to show quality spectrogram,* the default color range on this chart represents values from 0 to 100.

Real-Time FFT

The Real-time FFT chart displays the instantaneous Fast Fourier Transform (FFT) signature of the RF signal seen by the radio. The Fast Fourier Transform (FFT) converts an RF signal from time domain to frequency domain. The frequency domain representation divides RF signals into discrete frequency bins; small frequency ranges whose width depends on the resolution bandwidth of the spectrum monitor (that is, how many Hz are represented by a single signal strength value). Each frequency bin has a corresponding signal strength value. Because there may be a large number of FFT signatures received by the radio every second, an algorithm selects one FFT sample to display in the Real-time FFT chart every second.



A hybrid AP on a 20 MHz channel will see 40 MHz Wi-Fi data as non-Wi-Fi data.

This chart can show an average for all samples taken over the last second, the maximum FFT power measured for all samples taken over ten channel sweeps, and the greatest FFT power recorded since the chart was last reset. When you hover your mouse over any line, a tooltip shows the power level and channel or frequency level represented by that point in the graph. When you hover your mouse over a frequency level (within the blue brackets on the graph), a tooltip shows the types of devices seen on that frequency, BSSID, power level, channels affected, and the time the device was last seen by the spectrum monitor.

Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 174: *Real-Time FFT Options*

Parameter	Description
Band	Radio band displayed in this graph. For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5 GHz Low**, **5 GHz Center** or **5 GHz High** to display data for that portion of the 5GHz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80MHz** option for very-high-throughput channels.
X-Axis	Select either **Channel** or **Frequency** to show FFT power for a range of channels or frequencies. If you select **Frequency**, you must select the radio frequency on which this chart should center, and determine the span of frequencies for the graph.
Channel Range	If you selected **Channel** in the **X-Axis** parameter, you must also specify a channel range to determine which channels appear in the X-axis of this chart. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the chart. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.
Y-axis	Select the range of power levels, in -dBm, to appear in the y-axis of this chart. Enter the lower value in the right field, and the higher value in the left field.
Show	Select the check box by the following items to display that information on the FFT Power chart. ▪ **Average:** the average power level of all samples recorded during the last 10 sweeps. ▪ **Max** The highest power recorded during the last 10 channel sweeps. ▪ **Max Hold:** the highest maximum power level recorded since the chart data was reset. To clear this setting, click the down arrow at the end of the title bar for this graph and select **Clear Max Hold**.

Swept Spectrogram

A spectrogram is a chart that shows how the density of the quantity being plotted varies with time. The spectrum analysis Swept Spectrogram chart plots real-time FFT Maximums, real-time FFT Averages, or the FFT Duty Cycle. In this swept spectrogram, the x-axis represents frequency or channel and the y-axis represents time. Each line in the swept spectrogram corresponds to the data displayed in the Real-Time FFT or FFT Duty Cycle chart.



A hybrid AP on a 20 MHz channel will see 40 MHz Wi-Fi data as non-Wi-Fi data.

The power or duty cycle values recorded in each sweep are mapped to a range of colors. In the average or maximum FFT power Swept Spectrogram charts, the signal strength levels are indicated by a range of colors between dark blue, which represents -90 dBm, and red, which represents a higher -50 dBm. The duty cycle Swept Spectrogram chart shows the percentage of the time tick interval that the selected channel or frequency was broadcasting a signal. These percentages are indicated by a range of colors between dark blue, which represents a duty cycle of 0% percent, and red, which represents a duty cycle

of 100%.

A spectrogram plot is a complex chart that can display a lot of information. If you are not familiar with these types of charts, they may be difficult to interpret.

The spectrum analysis Swept Spectrogram measures FFT power levels or duty cycle data each second, so after every 1-second sweep, the newest data appears as a thin colored line on the bottom of the chart. Older data is pushed up higher on the chart until it reaches the top of the spectrogram and ages out.

Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 175: *Swept Spectrogram Options*

Parameter	Description
Band	Radio band displayed in this graph. For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5 GHz Low**, **5 GHz Center** or **5 GHz High** to display data for that portion of the 5GHz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80MHz** option for very-high-throughput channels.
X-Axis	Select either **Channel** or **Frequency** to show FFT power or duty cycles for a range of channels or frequencies. If you select **Frequency**, you must select the radio frequency on which this chart should center, and determine the span of frequencies for the graph.
Channel Range	If you selected **Channel** in the **X-Axis** parameter, you must also specify a channel range to determine which channels appear in the x-axis of this chart. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the chart. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.
Color-Map Range	*If this chart is configured to show average or maximum FFT values,* the default color range on this chart represents values from -50dBm (red) to -90dBm (blue). If you would like the color range on this chart to represent a different range of FFT power levels, enter this range in the **from** and **to** entry blanks. For example, if you defined a color-map range from -60 to -80, then any FFT power level at or above -60 dBm appears as red, and any FFT power level at or below -80 appears blue. Only the channel or frequency qualities between -60 dBm and -80 dBm would be represented by gradiented colors within the color range. *If this chart is configured to show the FFT duty cycle*, the default color range on this chart represents duty cycles from 0% (red) to 100% (blue). If you would like the color range on this chart to represent a different range of FFT duty cycle percentages, enter this range in the **from** and **to** entry blanks. For example, if you defined a color-map range from 25 to 75, then any FFT duty cycle at or below 25% appears as red, and any FFT duty cycle at or below 75% appears blue. Only the duty cycle levels between 25% and 75% would be represented by gradiented colors within the color range.

Parameter	Description
	NOTE: If your swept spectrogram is showing a single color only, you may need to increase the color map range to display a greater range of values.
Show	Select **FFT Avg**, **FFT Max** or **FFT Duty Cycle** to select the type of data you want to appear in this chart.

Working with Non-Wi-Fi Interferers

The following table describes each type of non-Wi-Fi interferer detected by the spectrum analysis feature. These devices appear in the following charts:

- Active Devices
- Active Devices Table
- Active Devices Trend
- Device Duty Cycle
- Device vs Channel
- Interference Power

Table 176: *Non-Wi-Fi Interferer Types*

Non-Wi-Fi Interferer	Description
Bluetooth	Any device that uses the Bluetooth protocol to communicate in the 2.4 GHz band is classified as a *Bluetooth* device. Bluetooth uses a frequency hopping protocol.
Fixed Frequency (Audio)	Some audio devices such as wireless speakers and microphones also use fixed frequency to continuously transmit audio. These devices are classified as *Fixed Frequency (Audio)*.
Fixed Frequency (Cordless Phones)	Some cordless phones use a fixed frequency to transmit data (much like the fixed frequency video devices). These devices are classified as *Fixed Frequency (Cordless Phones)*.
Fixed Frequency (Video)	Video transmitters that continuously transmit video on a single frequency are classified as *Fixed Frequency (Video*). These devices typically have close to a 100% duty cycle. These types of devices may be used for video surveillance, TV or other video distribution, and similar applications.
Fixed Frequency (Other)	All other fixed frequency devices that do not fall into one of the above categories are classified as *Fixed Frequency (Other*). Note that the RF signatures of the fixed frequency audio, video and cordless phone devices are very similar, and that some of these devices may be occasionally classified as Fixed Frequency (Other).
Frequency Hopper (Cordless Base)	Frequency hopping cordless phone base units transmit periodic beacon-like frames at all times. When the handsets are not transmitting (that is, no active phone calls), the cordless base is classified as *Frequency Hopper (Cordless Base*).

Non-Wi-Fi Interferer	Description
Frequency Hopper (Cordless Network)	When there is an active phone call and one or more handsets are part of the phone conversation, the device is classified as *Frequency Hopper (Cordless Network*). Cordless phones may operate in 2.4 GHz or 5 GHz bands. Some phones use both 2.4 GHz and 5 GHz bands (for example, 5 GHz for Base-to-handset and 2.4 GHz for Handset-to-base). These phones may be classified as unique Frequency Hopper devices on both bands.
Frequency Hopper (Xbox)	The Microsoft Xbox device uses a frequency hopping protocol in the 2.4 GHz band. These devices are classified as *Frequency Hopper (Xbox*).
Frequency Hopper (Other)	When the classifier detects a frequency hopper that does not fall into one of the above categories, it is classified as *Frequency Hopper (Other*). Some examples include IEEE 802.11 FHSS devices, game consoles, and cordless or hands-free devices that do not use one of the known cordless phone protocols.
Microwave	Common residential microwave ovens with a single magnetron are classified as a *Microwave*. These types of microwave ovens may be used in cafeterias, break rooms, dormitories and similar environments. Some industrial, healthcare or manufacturing environments may also have other equipment that behave like a microwave and may also be classified as a Microwave device.
Microwave (Inverter)	Some newer-model microwave ovens have inverter technology to control the power output and may have a duty cycle close to 100%. These microwave ovens are classified as *Microwave (Inverter)*. Dual-magnetron industrial microwave ovens with higher duty cycle may also be classified as Microwave (Inverter). As in the Microwave category described above, there may be other equipment that behave like inverter microwaves in some industrial, healthcare or manufacturing environments. Those devices may also be classified as Microwave (Inverter).
Generic Interferer	Any non-frequency hopping device that does not fall into one of the other categories described in this table is classified as a *Generic Interferer*. For example, a Microwave-like device that does not operate in the known operating frequencies used by the Microwave ovens may be classified as a Generic Interferer. Similarly, wide-band interfering devices may be classified as Generic Interferers.

Understanding Spectrum Analysis Session Log

The spectrum analysis **Session Log** tab displays times the spectrum monitors and hybrid APs connected to or disconnected from the spectrum client during the current browser session. This tab also shows changes in the scanning channel caused by changes to the 802.11A or 802.11G radio profile or automatic channel changes by the DFS or ARM features of the hybrid AP. The latest entry in the session log is also displayed in a footer at the bottom of the Spectrum Monitors and Spectrum Dashboard window. When you close the browser and end your spectrum analysis session, the session log is cleared.

Viewing Spectrum Analysis Data

You can use the command-line interface to view spectrum analysis data from any spectrum monitor, even if that spectrum monitor is currently sending data to the WebUI of the another spectrum monitor client.

Table 177 shows the commands that display spectrum analysis data in the CLI interface.

Table 177: *Spectrum Analysis CLI Commands*

Command	Description
`show ap spectrum ap-list`	Shows spectrum data seen by an access point that has been converted to a spectrum monitor.
`show ap spectrum channel-metrics`	Shows channel utilization information for a 802.11A or 802.11G radio band, as seen by a spectrum monitor
`show ap spectrum channel-summary`	Displays a summary of the 802.11A or 802.11G channels seen by a spectrum monitor.
`show ap spectrum client-list`	Shows details for Wi-Fi clients seen by a specified spectrum monitor.
`show ap spectrum debug`	Sub-commands under this command save spectrum analysis channel information to a file on the stand-alone controller.
`show ap spectrum device-duty-cycle`	Shows the current duty cycle for devices on all channels being monitored by the spectrum monitor radio.
`show ap spectrum device-history`	Displays spectrum analysis history for non-interfering devices.
`show ap spectrum device-list`	Shows summary table and channel information for non-Wi-Fi devices currently seen by the spectrum monitor.
`show ap spectrum device-log`	Shows a time log of add and delete events for non-Wi-Fi devices.
`show ap spectrum device-summary`	Shows the numbers of Wi-Fi and non-Wi-Fi device types on each channel monitored by a spectrum monitor.
`show ap spectrum interference-power`	Shows the interference power detected by a 802.11A or 802.11G radio on a spectrum monitor.
`show ap spectrum monitors`	Shows a list of APs currently configured as spectrum monitors.
`show ap spectrum technical-support`	Saves spectrum data for later analysis by your Aruba technical support representative.

Recording Spectrum Analysis Data

The spectrum analysis tool allows you to record up to 60 continuous minutes (or up to 10 Mb) of spectrum analysis data. By default, each spectrum analysis recording displays data for the Real-Time FFT, FFT Duty Cycle, Interference Power and Swept Spectrogram charts, however, you can view recorded device data for any the spectrum analysis charts supported by that spectrum monitor radio. Configurable recording settings allow you to start a recording session immediately, or schedule a recording to begin at a later date and time. Each recording can be scheduled to end after a selected amount of time has passed, or continue on until the recorded data file reaches a specified size. You can save the file to your spectrum monitor client, then play back that data at a later time.

The following sections provide information on creating, saving, and playing spectrum analysis data:

- Creating a Spectrum Analysis Record on page 930
- Saving Recording on page 930
- Playing Spectrum Analysis Recording on page 930

Creating a Spectrum Analysis Record

To record spectrum analysis data for later analysis, log in to the managed device obtained from the **Managed Network** node hierarchy in the Mobility Master (as described in **Obtaining Managed Device Node Details** section of Spectrum Analysis Tasks on page 907):

1. In the **Managed Device** node hierarchy, navigate to **Diagnostics** > **Tools** > **Spectrum Analysis**.
2. Click **Launch** to launch the Spectrum Analysis tool.
3. Click the **Spectrum Monitors** tab in the new window.
4. Click **Record** at the top of the window. The **New Recording** popup window appears.
5. Click the **Record From** link, and select the spectrum monitor whose data you want to record.
6. Next, decide whether you want the recording to start immediately, or at a later scheduled time. If you want the recording to start immediately, select **When the OK button is clicked**. To schedule a different starting time for the recording, click the date and time drop-down lists to select a starting month, day, year and time.
7. The recording continues until either the specified amount of time has passed, or until the recording files reaches a selected size. Click the **Length of recording reaches** drop-down list and select the amount of time the recording should last, or click the **Data file reaches** drop-down list and select the maximum file size for the recording.
8. Click **OK** to save your settings. If you selected the **When the OK button is clicked** in step 5, the recording begins.

While the recording is in progress, a round, red recording icon and recording status information appears at the top of the spectrum dashboard. You can view data for other spectrum monitors and charts while the recording is in progress. If you want to stop the recording before recording period has finished, click **Stop** by the recording status information. When you the **Stop**, a popup window appears and allows you to stop and delete the current recording, stop and save the recording in its current state (before it has completed), or continue recording again.

Saving Recording

After the recording has ended, either because the recording period has elapsed, the recording maximum file size has been reached, the **Spectrum Monitor Recording Complete** window appears and displays information for the current recording.

The following procedure saves the recording file:

1. From the **Spectrum Monitor Recording Complete** window, click **Continue**.
2. A **Save As** window appears and prompts you to select a file name for the recording and a location to save the file.
3. Click **Save**.

Playing Spectrum Analysis Recording

There are two ways to play back a spectrum recording. You can use the playback feature in the spectrum dashboard, or view recordings using the Aruba RFPlayback tool downloaded from the Aruba website.

Playing Recording in Spectrum Dashboard

The spectrum monitor does not have to be subscribed to your spectrum analysis client in order to play back a recording in the spectrum dashboard. However, you cannot play back an existing recording in the spectrum dashboard while another recording session is currently in progress.

To play a spectrum analysis recording in the spectrum dashboard, log in to the managed device obtained from the **Managed Network** node hierarchy in the Mobility Master (as described in **Obtaining Managed Device Node Details** section of Spectrum Analysis Tasks on page 907):

1. In the **Managed Device** node hierarchy, navigate to **Diagnostics** > **Tools** > **Spectrum Analysis**.
2. Click **Launch** to launch the Spectrum Analysis tool.
3. Click the **Spectrum Dashboards** tab.
4. Click **Load File for Playback**.
5. An **Open** dialog box appears and prompts you to browse to and select the file you want to open.
6. Click **Open**.
7. Click the triangular play icon at the top of the window to start playing back the recording.

Recorded data for the selected spectrum monitor and dashboard view appears in the spectrum analysis dashboard. You can replace any of the graphs in the playback window with a different graph type while replaying the recording. A playback progress bar at the top of the window shows what part of the recording currently appears on the dashboard. If you pause the recording, you can click and drag the red slider on this progress bar to advance to or replay any part of the current record.

Playing Recording Using RFPlayback Tool

The Aruba RFPlayback tool can play spectrum recordings created in this and earlier versions of ArubaOS. Aruba uses the Adobe AIR application to display spectrum recording information. If you have not done so already, follow the steps below to download and install the free Adobe AIR application and the Aruba spectrum playback tool.

1. Download the Adobe Air application from *http://get.adobe.com/air/* and install it on the client on which you want to play spectrum recordings.
2. Next, download the spectrum playback installation file from the Aruba website.
3. Open the folder containing the spectrum installation file, and double-click the spectrum.air icon to install the spectrum playback tool. You will be prompted to select the folder in which you want to install this tool.

Once you have installed the Aruba RFPlayback tool, follow the steps below to load and view a spectrum recording.

1. Start the Spectrum playback application.
2. Click **Load File for Playback**. An **Open** dialog box appears and prompts you to browse to and select the file you want to open.
3. Click the triangular play icon at the top of the window play the recording.

The RFPlayback tool also allows you to select and display different graph types while the recording playback is in progress. A playback progress bar at the top of the window shows what part of the recording is displayed in the playback tool. If you pause the recording, you can click and drag the red slider on this progress bar to advance to or replay any part of the current record.

The automatic reporting feature, also known as PhoneHome, allows a Mobility Master to send report events such as hardware failures, software malfunctions, and other critical events. The PhoneHome automatic reporting is disabled by default. When you enable the PhoneHome automatic reporting feature, the Mobility Master sends Aruba support weekly reports about the Mobility Master's configuration, licenses, software and hardware status, and any software malfunctions via Aruba Activate or a secure email. In the event that you need to contact Aruba support with a question about your Mobility Master, you can use this feature to generate and immediately send a status report, so that Aruba support can diagnose the issue with the most current Mobility Master data.

The PhoneHome feature can send reports to Aruba support through the Aruba Activate server using the HTTPS protocol (recommended for most deployments), or send reports to a local SMTP server in email messages.

For information to help you determine whether you should send PhoneHome reports to Aruba support via the Activate server or an SMTP server, see Registering with Activate on page 932

For procedures to configure this feature to send reports or to view your Mobility Master's report history, refer to the following topics:

- Configuring PhoneHome Automatic Reporting on page 933
- Viewing Report Status on page 935

Registering with Activate

Before sending PhoneHome reports using Activate, the managed device should be registered with the Activate server using the username or password authentication method.

The following CLI commands are used to register with Activate server:

```
(host) [md] (config) #activate
(host) [md] (activate) #username ztp
(host) [md] (activate) #password ztpadmin
(host) [md] (activate) #write memory
Saving Configuration...
Partial configuration for /mm/mynode
-------------------------------------
Contents of : /flash/config/partial/25/p=sc=mynode.cfg
activate
username "ztp"
password b8698637f6bc63bf5a851a16a2020b816907d92ba8b85a62
!
Configuration Saved.
(host) [md] (activate) #exit
(host) [md] (config) #exit
(host) [md] #show activate
(host) [md](config)# show activate
activate
--------
Parameter                                 Value                                 Set
---------                                 -----                                 ---
Activate Whitelist Service                Enabled
Activate URL                              https://activate.arubanetworks.com
```

```
Provision Activate URL                      https://device.arubanetworks.com
Activate Login Username                     ztp
Activate Login Password                     ********
Periodic Interval for WhiteList Download    1
Add-Only Operation                          Enabled
Custom cert to upload to Activate           N/A
Server cert to be used for IPSEC            N/A
```



The **Periodic Interval for WhiteList Download** parameter indicates the whitelist download period in days.

Configuring PhoneHome Automatic Reporting

Use the WebUI or the CLI to configure the Mobility Master to send weekly status reports. When you enable this feature, the Mobility Master sends reports every week by default.

The procedure to configure PhoneHome automatic reporting varies, depending upon whether you want to send reports via Aruba Activate or an SMTP email server. The following procedures describe the tasks to configure automatic reporting using Activate or SMTP.

Configuring PhoneHome Using Activate

The following procedure describes how to configure PhoneHome automatic reporting using Aruba Activate:

1. In the **Mobility Master** node hierarchy, select the device and navigate to **Diagnostics > Technical Support > TAC server** tab.
2. Toggle the **PhoneHome** switch to enable this feature.
3. Select **HTTPS** check box in the **HTTPS Configuration** field.
4. In the **Email-ID** field, enter a valid email address with a domain name associated with your Mobility Master. This field is used in the SMTP header and used to validate ownership with PhoneHome data.
5. Use the **Report Type** check boxes to specify whether you want automatic reporting to send regular weekly reports, or a single individual report.
   - Click **Auto Report** to schedule weekly status reports to be sent to Aruba.
   - Click **Report Now** to immediately send a single status report to Aruba.
6. Click **Apply**. If you selected the **Report Now** option in the **Report Type** field, the **Report Now** check box clears, indicating that no additional reports are scheduled to be sent.

Configuring PhoneHome Using SMTP

The following procedure describes how to configure PhoneHome automatic reporting to send reports through an SMTP email server:

1. In the **Mobility Master** node hierarchy, select the device and navigate to **Diagnostics > Technical Support > TAC server** tab.
2. Toggle the **PhoneHome** switch to enable this feature.
3. Select **SMTP** option in the **Protocol** field.
4. In the **Server IP Address** field, enter the IP address of your SMTP server. Optionally, in the **Server Port** field, enter the port the Mobility Master should use to access the server.

5. In the **Username** field, enter the user name of the user from whose email account the reports should be sent. Optionally, if your SMTP server requires the sender to be authenticated, enter a valid sender's user name and password in the **Username** and **Password** fields.
6. In the **Email-ID** field, enter the email address from which the reports should be sent. Optionally, if your SMTP server has limits on email attachment sizes, enter the attachment size in the **Max attachment size** field. Any status reports larger than this size is divided into multiple emails.
7. Use the **Report Type** check boxes to specify whether you want automatic reporting to send regular weekly reports, or a single individual report.
   - Click **Auto Report** to schedule weekly status reports to be sent to Aruba.
   - Click **Report Now** to immediately send a single status report to Aruba.
8. Click **Apply**.

You can disable automatic reporting at any time by returning to the **Diagnostics > Technical Support > TAC server tab** and either unchecking the **Auto Report** check box, or disabling the **PhoneHome** feature.

Configuring PhoneHome Using the CLI

The following CLI commands configure automatic reporting, and identifies the Aruba Activate or SMTP server that you want to use to send these messages:

```
(host) [mynode] (config) #phonehome
auto-report
https
smtp
```



Your SMTP and Activate server settings are preserved even when automatic reporting is disabled.

Sending Reports to Activate vs. SMTP Servers

By default, Mobility Master sends PhoneHome reports to the Activate server using HTTPS.

Most deployments should retain the default behavior introduced in this release and send PhoneHome reports using Activate. However, if the Mobility Master is behind a proxy server and does not have direct access to Internet, PhoneHome should be configured to send reports using SMTP. The following section describes the benefits of each of these configurations options.

Sending PhoneHome Reports Using Activate

PhoneHome integration with Activate offers the following benefits:

- **Simpler configuration:** PhoneHome only requires you to configure the email ID of the network administrator managing the device, as Activate already has information to accurately identify your Mobility Master. If a DNS server is not configured on the Mobility Master, PhoneHome will query the public DNS service (8.8.8.8) to resolve the Activate server IP address.
- **Smaller bandwidth requirements:** When the PhoneHome feature sends the report to the Activate server, the PhoneHome report is zipped into a smaller package, then divided into smaller 1 MB pieces before being sent to the server using secure HTTPS. Only reports sent to Activate are zipped before they are sent, so reports sent to Activate use less bandwidth than a report sent to a SMTP server.
- **Enhanced error management:** If any individual portion of the report is not successfully received by the Activate server, PhoneHome makes up to three attempts to resend just that portion of the file

rather than resending the entire report. Reports sent via SMTP must be resent in their entirety if any portion is not received by the SMTP server.

- **Automatic removal of old reports:** Once the entire report has been sent to the Activate server, Activate sends an acknowledgment to the Mobility Master, prompting the Mobility Master to delete its local copy of the report.
- The PhoneHome feature can be enabled or disabled using the **Diagnostics > Technical Support > TAC server**tab in the WebUI. The same can also be done through **phonehome [enable | disable]** option in the CLI.

Sending Reports Using SMTP

If you configure the PhoneHome feature to use SMTP, the PhoneHome status reports are sent via email. When the Mobility Master generates the report email with the PhoneHome data file attachment, it forwards the email to the local SMTP server configured on your local network, which then relays the message to Aruba technical support. If your email server requires the sender to be authenticated before message delivery, the Mobility Master can connect to the SMTP server by supplying the sender's user name and password.

When PhoneHome reports are sent using SMTP, the PhoneHome report attachment is encrypted before it is transmitted to the SMTP server. It is then decrypted by Aruba support when it is received. If the PhoneHome status report email is larger than the maximum email size supported by your SMTP server, the Mobility Master divides the PhoneHome attachment into smaller attachments and sends the report to Aruba in multiple emails. If any individual portion of the report is not successfully received by the SMTP server, PhoneHome resends the entire report.

Sending an Individual Report

If you are currently experiencing a problem and have contacted Aruba about the issue, Aruba technical support may ask you to generate and send an individual report, which describes the Mobility Master's current status, and reports any software or hardware errors. Once this report has been successfully uploaded, you may receive an email that contains a unique reference number you can use to track your recently opened ticket.



If you have not yet enabled automatic reporting feature or defined an SMTP server for this feature, follow steps 1-9 of the WebUI procedure described in Configuring PhoneHome Automatic Reporting on page 933.

The following procedure describes how to generate and send a PhoneHome status report:

1. In the **Mobility Master** node hierarchy, select the device and navigate to **Diagnostics > Technical Support > TAC server** tab.
2. Click **Report Now** check box.
3. Click **Apply.**

The following CLI command generates and sends a PhoneHome status report:

```
(host) [mynode] (config) #phonehome now
```

Viewing Report Status

Both the WebUI and CLI can show the status of the automatic reporting feature from the Mobility Master's last reset, including whether this feature is enabled, and the number of report messages that were successfully sent or failed to reach Activate or the SMTP server.

The following procedure describes how to view the report status:

1. In the **Mobility Master** node hierarchy, navigate to **Diagnostics > Technical Support > TAC server** tab.
2. In the **Protocol** field, select your protocol option and expand the **PhoneHome Status** accordion to view statistics for sent reports for that particular protocol. The **Statistics** and the **Transaction History** table are displayed. These tables contain the following information:

Table 178: *Automatic Reporting Statistics*

Report Statistic	Description
Success	Number of reports successfully sent to Activate or the SMTP server.
Failed	Number of reports that failed to reach Activate or the SMTP server after one or more retry attempts.
Retries	Number of times the Mobility Master attempted to retry sending a report to Activate or the SMTP server.
Manual PhoneHome	Number of reports generated by the Mobility Master because the **Report Now** setting was enabled.
Auto-report	Number of weekly reports generated by the Mobility Master because the **Auto Report** setting was enabled.

Table 179: *Transaction History*

Report Statistic	Description
TRANSACTION ID	An ID number for a specific report transaction. This transaction ID includes a timestamp showing when the transaction was first attempted.

The following CLI command displays statistics for Automatic Reporting settings and report status:

```
(host) [mynode] #show phonehome
global              Display Phonehome global settings
history             Display a history of phonehome transactions
report-status       Display status of reports uploaded to Aruba TAC Server stats PhoneHome Statistics
```

PhoneHome-Lite

PhoneHome-Lite is an HTTPS-based tracking tool used to monitor WebCC feature usage on each managed device. Aruba managed devices communicate with Activate servers over a secure HTTPS SSL through the PhoneHome infrastructure to send information about which users have enabled WebCC. This usage data can then be analyzed to determine the scope of future WebCC feature licensing.

The following procedure describes how to configure PhoneHome-Lite:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Services > Firewall >Global Setting**.

2. In the **Global Setting** accordion, toggle the **Enable Web Content Classification** switch to enable the feature.
3. On enabling, WebCC, the WebCC usage information will be sent to Aruba at every 7 days interval.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command enables PhoneHome-Lite:

```
(host) [md] #firewall web-cc
```

The following CLI command displays the WebCC configuration:

```
(host) [md] #show firewall
Global firewall policies
------------------------
Policy                                        Action          Rate       Port
------                                        ------          ----       ----
Enforce TCP handshake before allowing data    Disabled
Prohibit RST replay attack                    Disabled
.....
Web Content Classification            Enabled
....
```

Chapter 38
Management Access

This chapter describes management access and tasks for a user-centric network and includes the following topics:

- Configuring Certificate Authentication for WebUI Access on page 938
- Secure Shell on page 939
- Enabling RADIUS Server Authentication on page 942
- Connecting to AirWave Server on page 949
- Advanced Monitoring on page 950
- Clarity on page 952
- Custom Certificate Support for Remote AP on page 953
- Implementing Specific Management Password Policy on page 955
- Configuring Centralized Image Upgrades on page 957
- Managing Certificates on page 972
- Configuring SNMP on page 983
- Enabling Capacity Alerts on page 985
- Configuring Logging on page 986
- Enabling Guest Provisioning on page 989
- Managing Files on Managed Device on page 1004
- SCP Server Support on page 1008
- Setting System Clock on page 1010
- ClearPass Policy Manager Profiling with IF-MAP on page 1013
- Whitelist Synchronization on page 1014
- Downloadable Regulatory Table on page 1014
- Configuring Concurrent Sessions on page 1017
- Implementing Management User Audits on page 1017
- Implementing Password Validation on page 1017
- Maintaining Standard Mandatory Notice and Consent Banner on page 1018
- Zeroizing TPM Keys on page 1018

Configuring Certificate Authentication for WebUI Access

The managed device supports client certificate authentication for users accessing the WebUI. (The default is for username and password authentication.) You can use client certificate authentication only or client certificate authentication with username and password (if certificate authentication fails, the user can log in with a configured username and password).



Each managed device can support a maximum of ten management users.

To use client certificate authentication, you must do the following:

1. Obtain a client certificate and import the certificate into the managed device. Obtaining and importing a client certificate is described in Managing Certificates on page 972.
2. Configure certificate authentication for WebUI management. You can optionally also select username and password authentication.
3. Configure a user with a management role. Specify the client certificate for authentication of the user.

The following procedure describes how to configure certificate authentication:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Admin** tab and expand the **Admin Authentication Options** accordion.
2. Under **WebUI Authentication**, set **Client Certificate** to **Enabled**. You can select **Username/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.



By default, the **default-self-signed** certificate is used as the server certificate. For more details on **default-self-signed** certificate, see Managing Certificates on page 972.

1. Click **Submit**.
2. Click **Pending Changes**.
3. In the **Pending Changes** window, select the check box and click **Deploy changes**.
4. To configure the management user, navigate to the **Configuration > System > Admin** tab and expand the **Management User** accordion.
   a. Select **Enable Local Authentication** as needed.
   b. Click **Show users with certificate authentication** and click +.
   c. Select **WebUI**.
   d. Enter the username.
   e. Select the user role assigned to the user upon validation of the client certificate. Starting from ArubaOS 8.1.0.0, a new management role, **standard** role, is supported. This role has root privileges but cannot make changes to the management users.
   f. Enter the serial number for the client certificate.
   g. Select the name of the CA that issued the client certificate.
   h. Click **Submit**.
   i. Click **Pending Changes**.
   j. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure certificate authentication:

```
(host) [md] (config) #web-server profile
(host) [md] (Web Server Configuration) #mgmt-auth certificate
(host) [md] (Web Server Configuration) #switch-cert <certificate>
(host) [md] (Web Server Configuration) #!
(host) [md] (config) #mgmt-user webui-cacert <certificate-name> serial <number>
<username> <rolename>
```

Secure Shell

SSH is enabled by default in ArubaOS, and thus lets you log in using a username and password. You can enable SSH login by using public key authentication while leaving username and password authentication enabled, or you may disable the username and password authentication and leave only the public key authentication enabled. In the FIPS mode of operation, SSH is pre-configured to only use Diffie-Hellman Group 14 with AES-CBC-128, AES-CBC-256, HMAC-SHA1, or HMAC-SHA1-96. These settings are not configurable.

When you import an X.509 client certificate into the managed device, the certificate is converted to SSH-RSA keys. When you enable public key authentication for SSH, the managed device validates the credentials of the client with the imported public keys. You can specify public key authentication only, or public key authentication with username and password (if the public key authentication fails, the user can login with a configured username and password).

Enabling Public Key Authentication

The managed device allows public key authentication of users accessing the managed device using SSH. (The default is for username and password authentication.)

To use public key authentication, you must do the following:

1. Import the X.509 client certificate into the managed device using the WebUI, as described in Managing Certificates on page 972
2. Configure SSH for client public key authentication. You can optionally also select username and password authentication.
3. Configure the username, role and client certificate.

The following procedure describes how to enable public key authentication.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Admin** tab and expand the **Admin Authentication Options** accordion.
2. Under **SSH (Secure Shell) Authentication Method**, set **Client Public Key** to Enabled. You can optionally select **Username/Password** to use both username and password and public key authentication for SSH access.
3. Click **Submit**.
4. To configure the user, navigate to the **Configuration > System > Admin** tab.
   a. Expand the **Management User** accordion.
   b. Click **Show users with certificate authentication**.
   c. Click **+**.
   d. Select **CLI through SSH** from **Interface to connect** drop-down list.



ArubaOS recommends that the username and role for SSH be the same as for the WebUI Certificate. You can optionally use the check box to copy the username and role from the Web Certificate section to the SSH Public Key section.

   e. Enter the **User name**.
   f. Select the management role assigned to the user upon validation of the client certificate.
   g. Select the **Client certificate**.
5. Click **Submit**.

The following CLI commands enable public key authentication.

```
ssh mgmt-auth public-key [username/password]
mgmt-user ssh-pubkey client-cert <certificate> <username> <role>
```

Enabling Ciphers and MAC Algorithms

You can configure SSH to enable or disable the following ciphers and MAC algorithms based on your preference:

- **AES-CBC**
- **AES-CTR**
- **HMAC-SHA1**
- **HMAC-SHA1-96**
- **HMAC-SHA2-256**

By default, all the algorithms are enabled. However, the managed device allows you to enable or disable a specific cipher or the HMAC-SHA1-96 authentication algorithm.

The following procedure describes how to enable a cipher encryption:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Admin** tab and expand the **Admin Authentication Options** accordion.
2. Under **SSH (Secure Shell) Authentication Method > Encryption**, select **AES-CBC**, **AES-CTR**, or **Both**.
3. Click **Submit**.

The following procedure describes how to enable HMAC-SHA1-96 authentication:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Admin** tab and expand the **Admin Authentication Options** accordion.
2. Under **SSH (Secure Shell) Authentication Method > Authentication**, select **HMAC-SHA1-96**.
3. Click **Submit**.

The following procedure describes how to enable HMAC-SHA2-256 authentication:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Admin** tab and expand the **Admin Authentication Options** accordion.
2. Under **SSH (Secure Shell) Authentication Method > Authentication**, deselect **HMAC-SHA1** and **HMAC-SHA1-96**.
3. Click **Submit**.

The following CLI command enables AES-CBC on the SSH server.

```
(host) [md] (config) #ssh disable-ciphers aes-ctr
```

The following CLI command enables AES-CTR on the SSH server.

```
(host) [md] (config) #ssh disable-ciphers aes-cbc
```

The following CLI command enables both the ciphers on the SSH server.

```
(host) [md] (config) #no ssh disable-ciphers
```

The following CLI command enables HMAC-SHA1.

```
(host) [md] (config) #ssh disable-mac hmac-sha1-96
```

The following CLI command enables both the MAC authentication algorithms on the SSH server.

```
(host) [md] (config) #no ssh disable-mac
```

Viewing Cipher and MAC configuration

The following CLI command shows the status of cipher and MAC configuration.

```
show ssh
```

Disabling Console Access

A new command is introduced to disable the console-login. The purpose of this command is to introduce an ability to lock down all console ports, for example, micro USB, mini USB on the managed device to enable high level security.



With this command, only console access over serial port, USB, and mini USB will be blocked. SSH or telnet is still allowed.

The following CLI command disables the console.

```
(host) [mynode] (config) #mgmt-user console-block
PLEASE SAVE THE CONFIGURATION. CONSOLE WILL BE BLOCKED ONCE USER LOGS OUT FROM
SERIALCONSOLE.
```

The following CLI command enables the console.

```
(host) [mynode] (config) #no mgmt-user console-block
```

Enabling RADIUS Server Authentication

This section contains the following topics that describe the different types of RADIUS server configuration and related procedures:

- Configuring RADIUS Server Username and Password Authentication on page 942
- Configuring RADIUS Server Authentication with VSA on page 943
- Configuring RADIUS Server Authentication with Server Derivation Rule on page 944
- Configuring Set-value Server-derivation Rule on page 945
- Disabling Authentication of Local Management User Accounts on page 946
- Verifying Configuration on page 947
- Resetting Admin Password on page 947
- Setting Administrator Session Timeout on page 949

Configuring RADIUS Server Username and Password Authentication

In this example, an external RADIUS server is used to authenticate management users. Upon authentication, users are assigned the default role root.

The following procedure describes how to configure the RADIUS server and the server group.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** page.
2. Configure the **RADIUS Server**.
   a. In the **All Servers** table, click **+**.
   b. Enter the **Name** for the server (for example, rad1).
   c. Enter the **IP address**.
   d. Select **Radius** as the **Type**.
   e. Click **Submit**.
   f. Click **Pending Changes**.
   g. In the **Pending Changes** window, select the check box and click **Deploy changes**.

3. Select **Server Groups** to display the **Server Group** list.
   a. Click + and enter the name of the new server group (for example, corp_rad).
   b. Click **Submit**.
   c. Click **Pending Changes**.
   d. In the **Pending Changes** window, select the check box and click **Deploy changes**.
4. Select the new server group created and click **+** in the **Server Group> <Server name>** table. A list of servers is displayed.
5. Select the new server from the list to map this server to the server group that was created. Click **Submit**.
6. Navigate to the **Configuration > System > Admin** page.
   a. Expand the **Admin Authentication Options** accordion.
   b. Under **Admin Authentication Options**, select a management role (for example, root) for the **Default Role**.
   c. Set **Enable** to **Enabled**.
   d. For **Server Group**, select the server group that you just configured.
   e. Click **Submit**.
   f. Click **Pending Changes**.
   g. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the RADIUS server and the server group.

```
aaa authentication-server radius rad1
  host <ipaddr>
  enable

aaa server-group corp_rad
  auth-server rad1

aaa authentication mgmt
  default-role root
  enable
  server-group corp_rad
```

Configuring RADIUS Server Authentication with VSA

In this scenario, an external RADIUS server authenticates management users and returns the Aruba VSA ID 4, Aruba-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. 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 managed device.

The Aruba VSA ID 42, Aruba-Admin-Path, can be used to specify a node in the Mobility Master hierarchy for which the administrative login is valid. A user will only be allowed to login to that node and its tree nodes. The full path of the node must be specified starting from /md for example, /md/company/country or /md/company/country/location/controller and it should also be noted that only a single node path can be specified. Access will be granted for the entire hierarchy starting at the specified node.



Users can still view the entire hierarchy regardless of the Aruba-Admin-Path specification.

Configuring RADIUS Server Authentication with Server Derivation Rule



Arubamanaged device does not make use of any returned attributes from a TACACS+ server.

A RADIUS server can return to the managed device 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 managed device. The value of the attribute can be either “root” or “network-operations” depending upon the user; the returned value is the role granted to the user.



Ensure that the value of the attribute returned by the RADIUS server is one of the predefined management roles. Otherwise, the management user will not be granted access to the managed device.

The following procedure describes how to configure RADIUS server authentication.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** page.
2. Configure the **RADIUS Server**.
   a. In the **All Servers** section, click +.
   b. Enter the **Name** for the server (for example, rad1).
   c. Enter the **IP address**.
   d. Select **Radius** as the **Type**.
   e. Click **Submit**.
   f. Click **Pending Changes**.
   g. In the **Pending Changes** window, select the check box and click **Deploy changes**.
3. The **Server Group** table displays the Server Group list.
   a. Click **+ in the Server Group** table and enter the name of the new server group (for example, corp_rad) and click **Submit**.
   b. Select the name to configure the server group.
   c. Click **+** in the **Server Group > <server name>** table and select a server from **Add an existing server** table.
   d. Select the server and navigate to **Server Group > <server name> > Server Rules**, click **+** to add a server rule.
   e. For Condition, select **Class** from the **Attribute** scrolling list. Select **value-of** from the **Operation** drop-down list. Select **Set Role** from the **Action** drop-down list.
   f. Click **Submit**.

g. Click **Pending Changes**.

h. In the **Pending Changes** window, select the check box and click **Deploy changes**.

4. Navigate to the **Configuration > System > Admin** page.

a. Under **Admin Authentication Options**, select a management role (for example, root) for the **Default Role**.

b. Set **Enable** to **Enabled**.

c. For **Server Group**, select the server group that you just configured.

d. Click **Submit**.

e. Click **Pending Changes**.

f. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure RADIUS server authentication.

```
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 above example, the RADIUS server returns the attribute Class to the managed device; 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 Set-value Server-derivation Rule

The following procedure describes how to configure the set-value for a server-derivation rule

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** page.
2. Select a RADIUS Server from the **All Servers** table.

a. To add a new RADIUS server, click **+** in the **All Servers** table and enter the name for the server (for example, rad1) and click **Submit**.

b. Select the name to configure server parameters, such as IP address. The **Mode** check box is enabled by default to activate the server.

c. Click **Submit**.

d. Click **Pending Changes**.

e. In the **Pending Changes** window, select the check box and click **Deploy changes**.

3. Select a server group from the **Server groups** table to display the Server Group list.

a. To add a new server group, click **+** and enter the name of the new server group (for example, corp_rad) and click **Submit**.

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 **Add existing server** and click **Submit**.

e. Under Server Rules, click **+** to add a server rule.

f. For Condition, select an attribute from the **Attribute** scrolling list. Select **equals** from the **Operation** drop-down list. Enter **it**. Select **Set Role** from the **Action** drop-down list. For **Role**, select **root** from the drop-down list.

g. Click **Submit**.

h. Click **Pending Changes**.

i. In the **Pending Changes** window, select the check box and click **Deploy changes**.

4. Navigate to the **Configuration > System > Admin** tab.

a. Expand the **Admin Authentication options** accordion, select a management role (for example, read-only) for the **Default Role.**

b. For **Server Group**, select the server group that you just configured.

c. Click **Submit**.

d. Click **Pending Changes**.

e. In the **Pending Changes** window, select the check box and click **Deploy changes**.

In the CLI

The following CLI commands configure the set-value.

```
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 212.

Disabling Authentication of Local Management User Accounts

You can disable authentication of management user accounts in local switches if the configured authentication servers (RADIUS or TACACS+) are 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 servers 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.

The following procedure describes how to disable authentication of local management user accounts.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Admin** tab.
2. Expand the **Management Users** accordion, make sure that the **Enable Local Authentication** toggle switch is disabled.
3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command disables authentication of local management user accounts.

```
no mgmt-user localauth
```

Verifying Configuration

The following CLI command verifies if authentication of local management user accounts is enabled or disabled.

```
show mgmt-user local-authentication-mode
```

Resetting Admin Password

This section describes how to reset the password for the default administrator user account (**admin**) on the managed device. Use this procedure if the administrator user account password is lost or forgotten.

1. Connect a local console to the serial port on the managed device.
2. From the console, login into the managed device as a password recovery user. For information, read Password Recovery user.
3. Enter configuration mode by typing in **configure terminal**.
4. To reset the administrator user account password, use the **mgmt-user admin root** command.
5. Enter a new password for this account and retype the same to confirm.
6. Exit from the configuration mode and the user mode.

If you have defined a management user password policy, make sure that the new password conforms to this policy. For details, see Implementing Specific Management Password Policy on page 955.

The following is an example of how to reset the admin password as a default password recovery user. If you have configured an alternate password recovery user, use its credentials to login to the controller. The commands in bold type are what you enter:

```
User: password
Password: forgetme!
(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
```

Password Recovery user

A password recovery user is a management user with root rights that is used to reset the admin password in the event of a lost or forgotten password. Starting with ArubaOS 8.4.0.0, a configurable alternate password recovery user can be created in addition to the default password recovery feature.



- Password recovery access using either the default password recovery user or the alternate password recovery user is allowed only through the serial console of a controller.
- Password recovery users can be configured only through SSH sessions and serial console sessions with a controller and not through WebUI.
- Aruba recommends to enable the default password recovery user before generating and sharing the tech-support logs or configuration files with customer support.
- It is recommended that either the default password recovery user is disabled or the alternate password recovery user is configured when setting up the network to ensure. This is to ensure that there are no vulnerabilities.

In the event of a lost/forgotten password, the administrator can login to the controller and reset the admin password as the default password recovery user using the username **password** and the password **forgetme!**. The default password recovery user is defined and is enabled by default . Disabling the Default password recovery user is recommended if the network uses a TACACS server to authenticate its management users.

To disable the default password recovery user, execute the following command in the configuration mode:

```
(host) (config) #password-recovery-disable
```

To enable the default password recovery user, execute the following command in the configuration mode:

```
(host) (config) #no password-recovery-disable
```

Starting with ArubaOS 8.4.0.0, an alternate password recovery user with a username and password can be created to reset the admin password. The alternate user's username can be 16 characters long and the password can be 32 characters long. Configuring the alternate password recovery user automatically disables the default password recovery user. Configuring the alternate password recovery user is highly recommended if the network is managed locally.



The alternate password recovery user will not be shown in the management user section of the WebUI. This user role cannot be configured through the WebUI.

To configure the alternate password recovery user, execute the following command in the configuration mode:

```
(host) (config) #password-recovery-user <username>
Password:******
Re-Type password:******
```

To disable the alternate password recovery user, execute the following command in the configuration mode:

```
(host) (config) #no password-recovery-user
```

The following is an example to configure the alternate password recovery user:

```
(host) #configure terminal
Enter Configuration commands, one per line. End with CNTL/Z
(host) (config) #password-recovery-user recadmin
Password:******
Re-Type password:******
```

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

Use the **show mgmt-user** command to view the configured management users and the status of the default password recovery user.

The following is an example of the show mgmt-user command with the default password recovery user enabled.

```
(host) #show mgmt-user
Default password recovery user: Enabled
Management User Table
---------------------
USER    PASSWD  ROLE   STATUS
----    ------  ----   ------
admin   *****   root   ACTIVE
```

The following is an example of the show mgmt-user command when the alternate password recovery user is configured.

```
(host) #show mgmt-user
Default password recovery user: Disabled
Management User Table
---------------------
USER       PASSWD  ROLE    STATUS
----       ------  ----    ------
admin      *****   root    ACTIVE
recadmin   *****   passR   ACTIVE
```

Setting Administrator Session Timeout

The following CLI commands define a timeout interval for a WebUI session.

```
(host) [md] (config) #web-server profile
(host) [md] (Web Server Configuration) #session-timeout <session-timeout>
```

In the above command, **<session-timeout>** can be any number of seconds from 30 to 3600, inclusive.

The following CLI commands define a timeout interval for a CLI session.

```
(host) [md] (config) #loginsession timeout <value>
```

In the above command, **<val>** can be any number of minutes from 5 to 60 or seconds from 1 to 3600, inclusive. You can also specify a timeout value of 0 to disable CLI session timeouts.

Connecting to AirWave Server

AirWave is a powerful and easy-to-use network operations system that manages wireless, wired and remote access networks, as well as wireless and wired infrastructures and a wide range of third-party manufacturers.

Managed devices can use the **Configuration > System > AirWave** section of WebUI to quickly and easily connect the managed device to an AirWave server. The following table lists the information you will need to connect a managed device to an AirWave server.

Table 180: *AirWave Wizard Checklist*

Information	Description
AirWave IP address	IP address of the server.
SNMP version	Specify if the managed device and AirWave server should communicate using SNMP v2 or SNMP v3. SNMP v3 communication between a managed device and an AirWave server use SHA authentication and AES encryption.

Information	Description
Community string	For SNMP v2, select the community string used to authenticate requests or enter the new community string using the option New community string from the drop-down list.
Username	For SNMP v3, enter the name of the SNMP user.

Advanced Monitoring

Advanced Monitoring or AMON is used in Aruba WLAN deployments for improved network management, monitoring and diagnostic capabilities. ArubaOS enables Mobility Master and managed devices to periodically send these AMON feeds which typically contain status or statistical messages about the network devices to the Network Management Systems such as AMP without being prompted.

Traditionally, AMP used to send polls to SNMP on the controller on a regular interval to gather stats as well as detect additions, deletions, or changes to clients and APs. The SNMP tables were also huge with limited data. The scaling issues are resolved with the introduction of AMON. AMON sends out detailed messages to AMP automatically.

AMON communications between ArubaOS and AirWave is sent over 8211/udp (same as PAPI). AMON gathers all the client data for AirWave to scale for the huge amounts of data being sent by the controller to AirWave as SNMP has scaling issues.

The following sections provide additional information on the AMON feature:

Salient Features

The following are the salient features of AMON feature:

- Spectrum information, like discovered non-wifi interfering devices is only transmitted to AirWave over AMON.
- Access Point Channel Utilization information like Channel Busy, Interference, Receiving and Transmission time is only transmitted to AirWave over AMON.
- Client statistics like goodput and client association rate (speed) is only transmitted over AMON.
- Clients associating and leaving the user table are updated to AirWave more frequently via AMON vs. every 5 or 10 minutes via SNMP polling.

AMON Message Size Changes

Data communication between managed devices and AirWave servers has shifted from the SNMP model to the faster, more reliable, and scalable AMON model. Though the SNMP model can still be used to communicate data, users generally encounter delayed AirWave updates and high CPU usage.

The AMON packet size has been capped at a default value of 1264 bytes to reduce the amount of fragmentation and message loss that typically occurs in larger packet sizes, which can force customers to fall back to the SNMP model. Message size has been capped at 1264 bytes to allow for the addition of AMON headers and PAPI, UDP, or IP headers. Each packet only contains one message to further reduce the amount of overall message loss, as the loss of even a single fragment can render an entire message invalid. To reduce fragmentation in case of IPv6 AMON packets, the size can be configured to 1152 bytes.

The AMON packet size can be modified using the following CLI command:

```
amon msg-buffer-size <msg-buffer-size>
```

With the additional message load due to the smaller packet size and 1:1 message to packet ratio, output has also been increased from 10 second intervals to 1 second intervals to distribute packets more evenly, helping maintain a more stable and less congested traffic flow.

Secure AMON Feed to AirWave

Starting from ArubaOS 8.1 new enhancements have been made to the AMON infrastructure that includes DTLS support both on IPv4 and IPv6. DTLS provides secure communication over unreliable sessions with minimal overhead and still provides real-time response.

ArubaOS 8.1 supports four modes of AMON transportation.

Table 181: *Modes of AMON Transport*

Transport or Secure Modes	Cleartext	Secure
IPv4	UDP	DTLS
IPv6	UDP IPv6	DTLS over IPv6

The following procedure describes how to enable a new AMON receiver.

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Certificates** tab.
2. Click **+** in the **Import Certificate** table.
3. Update the following fields in the New Certificate table:
   - **Certificate name**
   - **Certificate filename**
   - **Optional passphrase**
   - **Retype passphrase**
   - **Certificate format**
   - **Certificate type** (ServerCert and TrustedCA)
4. Click **Submit**.
5. In the **Managed Network** node hierarchy, navigate to **Configuration > System > More > General** page.
6. Click **+** in the **MON Receivers** table.
7. Provide the server name.
8. Select a profile list from the **Profile list** drop-down list.
9. Select **UDP** from the Transport type drop-down list.
10. Select the **Secure** check box.
11. Click **Submit**.

The following CLI command specifies the certificate name.

```
(host) [md] (config) #amon cert <cert-name>
```

Support for Smart AMON

Starting from ArubaOS 8.4.0.0, Cloud based monitoring of Campus controllers becomes very essential. So, the AMON feeds are now made programmable and cloud friendly to help minimize the AMON telemetry traffic between the controller and the Cloud.

This feature enables the following new capabilities to the AMON feeds to Cloud destination over Websockets.

- Per-destination-tuned low-over head AMON feed
  - Programmable AMON feed with new controller APIs
  - Dynamic and precise selection of messages for cloud consumption
  - Flexible message size per destination
  - Adjustment of rate for selected AMON messages
- Cloud bootstrapping support
  - Cloud-optimized to minimize latency and overhead (including compression)
  - Allows quick state reconciliation after websocket connection or re-connection
  - Bulk micro-bootstrapping support (i.e. request for specific number of objects in one request)
  - AMON feed compression
- To reduce bandwidth consumption towards cloud

Clarity

The Aruba Clarity feature in AirWave proactively analyzes end-users' quality of experience by providing enhanced monitoring capabilities for critical network services, such as time and response failures for a mobile device to associate with a Wi-Fi radio. Other services monitored include authentication time period to a RADIUS server, gathering an IP address through DHCP, and resolution of names for DNS services. This gives IT organizations end-to-end visibility into problems before they escalate as metrics are monitored in real-time, and also captured through on-demand or scheduled testing for predictive insight.

The following sections provide information on Clarity Live and Clarity Synthetic features.

Clarity Live

Clarity Live helps diagnose client connectivity issues. It provides the network administrator or engineers with more information regarding the exact stage at which the client connectivity fails or provides data where the DHCP or RADIUS server is slow.

The managed device collects all information related to user transitions such as association, authentication, and DHCP. Then, the managed device sends these records to a management server such as AirWave. The management server analyzes the data and concludes which DHCP or RADIUS server was not working efficiently causing user connectivity issues. This enhancement allows the management server to isolate WLAN issues caused by external servers such as DHCP or RADIUS.



In this release, the Clarity Live data is not displayed or available for bridge mode Remote AP users.

Following are the advantages of Clarity Live:

- Improves user serviceability.
- Provides network administrator and engineers information on the client connectivity failures.
- Easier DHCP debugging.

The following topic describes how to enable Clarity Live Statistics:

Enabling Clarity Live Statistics

You can enable the Clarity Live statistics using the **mgmt-server profile** command in the CLI interface.

Execute the following command to enable the Clarity Live statistics from the AP:

```
(host) [md] (config) #mgmt-server
(host) [md] (config) #mgmt-server profile default-amp
```

```
(host) [md] (Mgmt Config profile "default-amp") #inline-ap-stats
```

Execute the following command to enable Clarity Live statistics related to authentication:

```
(host) [md] (Mgmt Config profile "default-amp") #inline-auth-stats
```

Execute the following command to enable Clarity Live statistics of DHCP

```
(host) [md] (Mgmt Config profile "default-amp") #inline-dhcp-stats
```

Execute the following command to enable Clarity Live statistics of DNS

```
(host) [md] (Mgmt Config profile "default-amp") #inline-dns-stats
```

Clarity Synthetic

Mobility Master provides support for Clarity Synthetic, which helps in detecting network health by using synthetic transaction from a Wi-Fi client. This feature converts the radios of a supported access point to switch from AP mode to station mode. The managed device converts one or both of the radios of the AP to station mode based on the instruction from a NMS. When the radio of the AP is in station mode, it starts synthetic data transaction within the network.

The following table provides a list of AP platforms that support Clarity Synthetic:

The network health is determined based on the response from the network and the time taken for the synthetic data transaction. The results captured as part of these transactions are used for the following purposes:

- Troubleshoot a live network
- Provide whole network overview (WLAN and Wired)
- Support Wi-Fi and Internet Protocol Service Level Agreement
- Troubleshoot Remote network using client traffic (Synthetic)

Custom Certificate Support for Remote AP

As Suite-B mandates using the AES-GCM encryption and ECDSA certificates for security, this feature allows you to upload custom RSA and ECDSA certificates to a Remote AP. This allows custom certificates to be used for IKEv2 negotiation which establishes a tunnel between the Remote AP and the managed device. Feature support includes the ability to:

- Upload a single CA certificate and Remote AP certificate which have either elliptical crypto key parameters with ECDSA or RSA parameters for signing and verification.
- Store the certificate in the flash of the Remote AP
- Store CSR and private key files in a USB
- Delete certificates
- Generate a CSR paired with a private key generation for the Remote AP. The private key is stored in the flash and the CSR can be exported out of the Remote AP to get it signed by the CA.

If there is a custom certificate present in the flash when rebooting, this feature creates a suite B tunnel with the managed device if the certificates uploaded are using EC algorithms. Otherwise it creates a tunnel using standard Remote AP IPsec parameters.

Suite-B Support for ECDSA Certificate

If a custom ECDSA certificate is present in the flash of a certificate-based Remote AP, it is automatically designated as a Suite-B Remote AP. On the managed device side, tunnel creation uses the server certificate as a default VPN server certificate.

Administering Suite-B support for a Remote AP includes these steps which are described in the following sections:

1. Setting the Default Server Certificate
2. Importing a custom certificate
3. Generating a CSR
4. Uploading the certificate

Setting Default Server Certificate

Enter the following commands to set the default server certificate that is presented to the Remote AP as the default VPN server certificate:

```
(host) [md] (config) #crypto-local isakmp server-certificate
<server_certificate_name>
```

Enter the following commands to add the CA certificate to verify the Remote AP certificate:

```
(host) [md] (config) #crypto-local isakmp ca-certificate <trusted CA>
```

Importing a Custom Certificate

Certificates can only be imported to the managed device using the WebUI.

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Certificates** and upload the certificate.
2. To use imported certificates to create a tunnel, in the Managed Device node hierarchy, navigate to **Configuration > Services > VPN > Certificates for VPN Clients**.

Generating CSR

The Remote AP console page allows you to generate a CSR. This is done through a private key which can be generated and saved to the Remote AP flash. A corresponding CSR is exported so it can be signed by the required CA to use as the Remote AP certificate. This Remote AP certificate can then be uploaded using the Upload button on the Remote AP Console page.

The subject of the Remote AP certificate needs to be the MAC address of the Remote AP, and nothing more. Note that this is case insensitive.

If you create a CSR on the Remote AP and then have a certificate issued by a CA, you must have the certificate in PEM format before uploading it to the Remote AP.

Uploading Certificate



When using the “rapconsole.arubanetworks.com” page on a bridge or split-tunnel Remote AP to manage certificates on the Remote AP, a blank page or a page that does not have the Certificates tabs on it may display. The Remote AP provisioning page that is standard on the Remote AP may conflict with the “rapconsole” page and thus confuse the browser. If this occurs, clear your browser cache first or use two different browsers.

The Upload button on the Remote AP console page that lets you upload the certificates to the Remote AP flash. The certificate needs to be in PEM format and uploading the Remote AP certificate requires that the corresponding private key is present in the Remote AP flash. Or, use the PKCS12 bundle where the chain includes the Remote AP private key with the Remote AP and CA certificates are optionally password protected.

Storing CSR and Private Key Files in USB

To provision a Remote AP to store the CSR and private key in a USB, use one of the following options:

AP Boot Prompt

At the AP boot prompt, issue the **setenv usb_csr 1** and **setenv usb_type 100** commands.



If this option is used to provision the Remote AP to store the files in the USB device, after the files are saved in the USB, enter the AP boot prompt to issue the **setenv usb_csr 0** command. This is mandatory.

The following procedure describes how to store CSR and private key files in USB.

1. In the **Managed Network** node hierarchy, navigate to **Configuration** > **Wireless** > **AP Installation** > **Provisioning**.
2. Select the **RAP**, click **Provision**.
3. Under **USB Settings**, select the **USB Parameters** check box.
4. Select the **USB storage for CSR/Key** check box.
5. Select **Device Type** as **storage**.
6. Click **Apply and Reboot**.

The following CLI commands store CSR and private key files in USB.

```
(host) [md] (config) #provision-ap
(host) [md] (AP provisioning) #read-bootinfo ap-name <ap name>
(host) [md] (AP provisioning) #usb-csr
(host) [md] (AP provisioning) #usb-type storage
```

Remote AP Console

The following procedure describes how to store CSR and private key files in USB using the Remote AP console.

1. In the **Managed Network** node hierarchy, navigate to **Configuration** > **Management** > **Certificates**.
2. For **Store CSR and key in USB/Flash**, select **USB** from the drop-down list.

After the Remote AP is provisioned to store the CSR and private key in a USB, log in to the Remote AP console, export the CSR and private key files to the USB. A **.p12** certificate file format must be manually created as the Remote AP certificate in the USB to bring up the IKE or IPsec connection.

Implementing Specific Management Password Policy

By default, the password for a new management user has no requirements other than a minimum length of 6 alphanumeric or special characters. However, if your company enforces a best practices password policy for management users with root access to network equipment, you may want to configure a password policy that sets requirements for management user passwords.

This section describes the following topics:

1. Defining Management Password Policy on page 955
2. Management Authentication Profile Parameters on page 957

Defining Management Password Policy

The following procedure describes how to configure specific management password policy setting.

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. Expand **Other Profiles**.

3. Select **Mgmt Password Policy**.
4. Configure the settings described in Table 182.

Table 182: *Management Password Policy Settings*

Parameter	Description
Enable Password Policy	Select this check box to enable the password management policy. The password policy will not be enforced until this check box 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 (!, @, #, $, %, ^, &, *, <, >, {, }, [,], :, ., comma, |, +, ~, `)	The minimum number of special characters. Range: 0-10 characters.
Username or Reverse of username NOT in Password	When you select this check box, the password cannot be the current username or the username spelled backwards of the management users.
Maximum consecutive character repeats	The maximum number of consecutive repeating characters allowed in a management user password. Range: 0-10 characters. By default, there is no limitation on the numbers of character that can repeat within a password, and the parameter has a default value of 0 characters.
Maximum Number of failed attempts in 3 minute window to lockout user	The number of failed attempts within a 3 minute window that causes the user to be locked out for the period of time specified by the **Time duration to lockout the user upon crossing the "lock-out" threshold** parameter. Range: 0-10 attempts. By default, the password lockout feature is disabled, and the default value of this parameter is 0 attempts.
Time duration to lock out the user upon crossing the "lock-out" threshold	The duration in time that locks out the user upon crossing the lock out threshold. Range: 0-60 in minutes.

5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures specific management password policy settings:

```
aaa password-policy mgmt
```

Management Authentication Profile Parameters

Table 183 describes configuration parameters on the Management Authentication profile page.



In the CLI, you configure these options with the **aaa authentication mgmt** and **aaa-server-group** commands.

Table 183: *Management Authentication Profile Parameters*

Parameter	Description
Enable	Enables authentication for administrative users.
Default Role	Select a predefined management role to assign to authenticated administrative users:
Root	Default superuser role
Guest-provisioning	Guest provisioning role
Location-api-mgmt	Location API role
Network-operations	Network operations role
No-access	No commands are accessible for this role
Read-only	Read-only role
No access	Negates any configured parameter.
Server Group	Name of the group of servers used to authenticate administrative users. See the CLI command **aaa-server-group**, in the *CLI Command Reference Guide* for more information.

Configuring Centralized Image Upgrades

The Centralized Image Upgrade feature uses ArubaOS images to upgrade the managed devices with the ArubaOS images hosted on an image server. When an upgrade action command is executed on the Mobility Master, the **upgrademgr** process running on Mobility Master sends an upgrade request to **upgrademgr** process running on corresponding managed devices. The managed devices then connect to the image server, validates the image file and downloads the appropriate image file. It then upgrades to the downloaded image file.

Starting from ArubaOS 8.2.0.0, IPv6 address support is added for Centralized Image Upgrade.

ArubaOS now provides native IPv6 support that allows the **upgrademgr** process to send upgrade requests between Mobility Master and managed devices in a native IPv6 deployment. If the Mobility Master IP address contains both IPv4 and IPv6 addresses, then IPv6 address is preferred over IPv4 address.

This section describes the following topics:

- Rebooting Mobility Master
- Configuring Boot Parameters on page 959
- Configuring Service Module Packages on page 959

- Configuring Upgrade Profile of Managed Device on page 961
- Configuring Upgrade Profile of Mobility Master on page 962
- Scheduling Upgrade of Managed Devices and Clusters on page 963
- Backing Up the Flash File System on page 969
- Copying the Flash File System on page 970
- Backing Up the Configuration on page 970
- Restoring the Flash File System on page 970
- Restoring the Configuration on page 970
- Clearing the Configuration on page 970
- Synchronizing the Database on page 971
- Exporting the WMS Database on page 971
- Importing the WMS Database on page 971
- Clearing Old Entries in the WMS Database on page 971
- Re-initializing the WMS Database on page 971
- Rebooting an AP on page 972

Configuring Upgrade Profile of Mobility Master

The following procedure describes how to configure an upgrade profile for the Mobility Master.

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Software Management > Upgrade**.
2. Configure the parameters described in the following table:

Table 184: *Upgrade Image Parameters for Mobility Master*

Parameter	Description
Upgrade using	Select the protocol used to download the image file. ■ FTP ■ Local File ■ SCP ■ TFTP ■ USB
Server IP address	Specify the IPv4 or IPv6 address of the image server. **NOTE:** Specify this parameter if **SCP**, **FTP**, or **TFTP** is used as the protocol to download the image file.
Image file name	Specify the image file name. **NOTE:** Use the **Browse** button to navigate to the image file if **Local File** is used as the protocol to download the image file.
Username	Specify the username to log in to the image server. **NOTE:** Specify this parameter if **FTP** or **SCP** is used as the protocol to download the image file.
Password	Specify the password to log in to the image server.

Parameter	Description
	NOTE: Specify this parameter if **FTP** or **SCP** is used as the protocol to download the image file.
Partition to upgrade	Select the partition to upgrade.
Reboot controller after upgrade	Enable the toggle switch to reboot the Mobility Master after upgrade.
Save current configuration	Select the check box if the existing configuration on the Mobility Master should be saved before rebooting. **NOTE:** The existing configuration on the Mobility Master is saved before it is rebooted only if the **Reboot controller after upgrade** toggle switch is enabled.

3. Click **Upgrade**.

The following CLI commands configure the upgrade profile on the Mobility Master in a native IPv6 deployment:

```
(host) [mynode] (config) #upgrade-profile
(host) [md] (Upgrade Profile) #serveraddr <serveraddr> username <username> password
<password> filepath <filepath> protocol scp
```

Rebooting Mobility Master

The following procedure describes how to reboot the Mobility Master.

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Software Management > Reboot**.
2. Select the **Save configuration before reboot** check box to save the existing configuration on the Mobility Master before rebooting.
3. Click **Reboot**.

Configuring Boot Parameters

The following procedure describes how to configure the boot parameters.

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Software Management > Boot Parameters**.
2. Select the boot partition (partition 0 or partition 1) to boot from.
3. Click **Apply**.

Configuring Service Module Packages

The following topics describe how to configure the service module packages on the Mobility Master:

- Viewing Service Module Packages on page 960
- Adding Service Module Package on page 960
- Activating Service Module Package on page 961
- Removing Service Module Package on page 961

Viewing Service Module Packages

You can view the list of service module packages using the WebUI or CLI.

To view the list of service module packages using the WebUI, in the **Mobility Master** node-hierarchy, navigate to **Maintenance > Software Management > Service Module Packages**.

To view the list of service module packages using the CLI, execute the following command:

```
(host) [mynode] #show packages
```

Adding Service Module Package

The following procedure describes how to add a service module package.

1. Obtain the required service package from the Aruba Support site.
2. In the **Mobility Master** node-hierarchy, navigate to **Maintenance > Software Management > Service Module Packages**.
3. Click **+** at the bottom of the **Service Module Packages** table.
4. In the **Add package** window, configure the parameters described in the following table:

Table 185: *Load New Package Configuration Parameters*

Parameter	Description
Access method	Select the protocol to access the service module package. ■ FTP ■ Local file ■ SCP ■ TFTP ■ USB
Host IP address	Specify the IPv4 or IPv6 address of the image server where the service module package resides. **NOTE:** Specify this parameter if **FTP**, **SCP**, or **TFTP** is used as the access method to access the service module package.
Image file name	Enter the exact service package name as residing on the image server. **NOTE:** Use the **Browse** button to navigate to the image file if **Local File** is used as the access method to access the service module package.
Destination file name	Specify the destination service module package name. **NOTE:** As a best practice, keep the image name same as destination file name.
Username	Specify the username to log in to the image server. **NOTE:** Specify this parameter if **FTP** or **SCP** is used as the access method.
Password	Specify the username to log in to the image server. **NOTE:** Specify this parameter if **FTP** or **SCP** is used as the access method.

5. Click **Submit** to validate the package.

The following CLI commands add a service module package.

```
(host) [mynode] #upgrade-pkg copy ftp: <ftphost> <username> <filename> flash:
<destfilename>
(host) [mynode] #upgrade-pkg copy scp: <scphost> <username> <filename> flash:
<destfilename>
(host) [mynode] #upgrade-pkg copy tftp: <tftphost> <filename> flash: <destfilename>
```

After downloading a service module package, the LSM feature performs the following compatibility checks to determine if the service module package is compatible with the running version of ArubaOS.

- **Platform Check**: Determines if the package must run on a specific platform.
- **Version Check**: Determines if the package version matches the version of ArubaOS running on the system.

If validation is successful, the installation process can continue. If validation is unsuccessful, the package is removed, and an error message appears.

Activating Service Module Package

The following procedure describes how to activate a service module package

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Software Management > Service Module Packages**.
2. Select a service module package from the **Service Module Packages** table.
3. In the **Service Module Packages > [name]** section, set **Status** to **Active**.
4. Click **Submit**.

The following CLI command activates a service module package.

```
(host) [mynode] #upgrade-pkg activate <packagename>
```

The service module package is halted and upgraded with the new service module package, during which time the service module package is unavailable to all users. After the service module package is activated, the service restarts.

Removing Service Module Package

The following procedure describes how to remove a service module package.

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Software Management > Service Module Packages**.
2. Select a service module package from the **Service Module Packages** table.
3. Click the **trash** icon.
4. Click **Delete** in the service module package delete window.

The following CLI command removes a service module package.

```
(host) [mynode] #upgrade-pkg remove <packagename>
```

Configuring Upgrade Profile of Managed Device

The following procedure describes how to configure an upgrade profile for a managed device:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. Under **All Profiles**, select **Controller Profile > Upgrade**.
3. Under **Upgrade Profile**, configure the parameters listed in the following table.

Table 186: *Upgrade Profile Parameters for Managed Device*

Parameter	Description
Server IP address	Specify the IPv4 address of the image server. This parameter is only used by managed devices running versions prior to ArubaOS 8.2 and accepts only IPv4 address. **NOTE:** For FTP or SCP protocol, specify the username and password.
Server IPv4/IPv6 address	Specify the IPv4 or IPv6 address of the image server. This parameter is only used by managed devices running version ArubaOS 8.2.0.0 or later versions. **NOTE:** For FTP or SCP protocol, specify the username and password.
User name	Specify the username of the image server.
Password	Specify the password of the image server.
Retype	Repeat the password of the image server.
Protocol	Select the protocol used to download the image file from the image server to the managed devices. ■ tftp ■ ftp ■ scp **NOTE:** Select none for local file.
File path	Specify the path of the image file on the image server.

4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

The following CLI command configures the upgrade-profile:

```
(host) [md] (config)#upgrade-profile
```

The following CLI command upgrades the managed device from the Mobility Master in a native IPv6 deployment:

```
(host) [mynode] #upgrade managed-devices copy-reboot configured-fileserver version <img-version>
```

The following CLI command initiates upgrade on all managed devices from the Mobility Master in a native IPv6 deployment:

```
(host) [mynode] #upgrade managed-devices copy-reboot fileserver scp <imagehost> <username> . version <img-version> all
```

Configuring Upgrade Profile of Mobility Master

The following procedure describes how to configure an upgrade profile for the Mobility Master.

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Software Management > Upgrade**.
2. Configure the parameters described in the following table:

Table 187: *Upgrade Image Parameters for Mobility Master*

Parameter	Description
Upgrade using	Select the protocol used to download the image file. ■ FTP ■ Local File ■ SCP ■ TFTP ■ USB
Server IP address	Specify the IPv4 or IPv6 address of the image server. **NOTE:** Specify this parameter if **SCP**, **FTP**, or **TFTP** is used as the protocol to download the image file.
Image file name	Specify the image file name. **NOTE:** Use the **Browse** button to navigate to the image file if **Local File** is used as the protocol to download the image file.
Username	Specify the username to log in to the image server. **NOTE:** Specify this parameter if **FTP** or **SCP** is used as the protocol to download the image file.
Password	Specify the password to log in to the image server. **NOTE:** Specify this parameter if **FTP** or **SCP** is used as the protocol to download the image file.
Partition to upgrade	Select the partition to upgrade.
Reboot controller after upgrade	Enable the toggle switch to reboot the Mobility Master after upgrade.
Save current configuration	Select the check box if the existing configuration on the Mobility Master should be saved before rebooting. **NOTE:** The existing configuration on the Mobility Master is saved before it is rebooted only if the **Reboot controller after upgrade** toggle switch is enabled.

3. Click **Upgrade**.

The following CLI commands configure the upgrade profile on the Mobility Master in a native IPv6 deployment:

```
(host) [mynode] (config) #upgrade-profile
(host) [md] (Upgrade Profile) #serveraddr <serveraddr> username <username> password
<password> filepath <filepath> protocol scp
```

Scheduling Upgrade of Managed Devices and Clusters

Starting with ArubaOS 8.4.0.0, you can schedule upgrade of managed devices and clusters.

This section provides information on the various scheduling upgrade methods:

- Scheduling Upgrade of Managed Devices on page 964
- Scheduling Upgrade of Managed Devices and Clusters on page 963

Scheduling Upgrade of Managed Devices

The following procedure describes how to schedule the upgrade of managed devices.

1. In the **Managed Network** node hierarchy, navigate to **Maintenance > Software Management**.
2. Select one or multiple managed devices from the table. The table lists the following columns:

Table 188: *Schedule Upgrade of Multiple Managed Devices*

Parameter	Description
NAME	Specifies the name of the managed device.
CURRENT VERSION	Specifies the ArubaOS release version (denoted by release number-build number) that is already running on the managed device. If an upgrade is in progress, this column displays the **Installation in progress** link. Hovering on the link opens a pop-up, which lists: ■ Additional information of the upgrade through the **Show details** link. ■ The time (in most significant non-zero units) when the upgrade was initiated. If an upgrade was initiated within the last 7 days, this column displays an upgrade status icon next to the value of the current version. Hovering on the upgrade status icon opens a pop-up, which lists: ■ Additional information of the upgrade through the **Show details** link. ■ The time (in most significant non-zero units) when the upgrade succeeded or failed. If an upgrade is scheduled, this column displays an upgrade schedule clock icon next to the value of the current version. Hovering on the upgrade schedule clock opens a pop-up, which lists: ■ The ArubaOS release version and scheduled date and time. ■ An option to reschedule the date and time. ■ An option to cancel the scheduled upgrade.
ACCESS POINTS	Specifies the total number of access points that are configured on the selected managed device.
GROUP	Specifies the parent group of the managed device.

3. In the **INSTALLATION SETTINGS > When** menu:
   - To upgrade the selected managed devices immediately, select **Now**.
   - To schedule the upgrade of the selected managed devices at a later date and time, select **Later**. Use the date and time picker to schedule the date and time. If the upgrade is scheduled to occur at a later date and time, a **Pending changes will automatically be deployed on managed controllers when the software installation will start** message is displayed.



- You can specify any date and time from the current date and time until the next 30 days.
- The upgrade occurs at the scheduled local date and time of the Mobility Master.

4. In the **Specify image file location, name and protocol to use for transfer > Server IP address** menu, specify the hostname, IPv4, or IPv6 address of the file server that hosts the ArubaOS image.
5. In the **Specify image file location, name and protocol to use for transfer > Image path** menu, specify the path of the ArubaOS image on the file server that hosts the image. You may use the **.** wildcard character to specify the default path.
6. In the **Specify image file location, name and protocol to use for transfer > Protocol** menu, select the protocol to transfer the ArubaOS image from the file server to the selected managed devices.

Table 189: *Protocol to Transfer ArubaOS Image*

Parameter	Description
FTP	Uses the FTP protocol to transfer the ArubaOS image from the file server to the selected managed devices. If you choose the FTP protocol, specify the username and password to log in to the FTP server.
SCP	Uses the SCP protocol to transfer the ArubaOS image from the file server to the selected managed devices. If you choose the SCP protocol, specify the username and password to log in to the SCP server.
TFTP	Uses the TFTP protocol to transfer the ArubaOS image from the file server to the selected managed devices.

7. In the **Specify image file location, name and protocol to use for transfer > Software to install** menu, specify the software release number of the ArubaOS image in release number-build number format.
8. In the **Specify image file location, name and protocol to use for transfer > Partition** menu, select the partition of the selected managed devices where the ArubaOS image should be installed.

Table 190: *Partition of Managed Device*

Parameter	Description
auto	Upgrades the default boot partition of the selected managed devices.
partition 0	Upgrades partition 0 of the selected managed devices.
partition 1	Upgrades partition 1 of the selected managed devices.

9. Click **Install**. A **Install Confirmation** pop-up displays the **Network performance will be temporary impacted during software installation Do you want to install software now?** message.
10. In the **Install Confirmation** pop-up, click **Install**.

To schedule the upgrade of a managed device using the CLI:

```
(host) [mynode] #upgrade managed-devices copy-reboot configured-fileserver version
8.4.0.0 all schedule 2018 12 25 12 30 00
```

To reschedule the scheduled upgrade of a managed device using the CLI:

```
(host) [mynode] #upgrade reschedule from 2018 11 30 23 59 59 to 2018 12 25 12 30 00 all
```

To cancel the scheduled upgrade of a managed device using the CLI:

```
(host) [mynode] #upgrade cancel-schedule 2018 12 01 23 59 59 all
```

Viewing Status of Scheduled Upgrade of Multiple Managed Devices

A scheduled upgrade icon in the banner indicates that one or more managed devices are scheduled for upgrade. Hovering over the scheduled upgrade icon opens a pop-up, which lists the number of managed devices that are scheduled for upgrade. Click the number in the pop-up to open the **Managed Network** node hierarchy table. The **Managed Network** node hierarchy table is filtered to list only the managed devices that are scheduled for upgrade.

If the upgrade of managed devices is in progress, a controller upgrade icon replaces the scheduled upgrade icon. Hovering over the controller upgrade icon opens a pop-up, which lists the number of managed devices that are currently under upgrade. Click the number in the pop-up to open the **Managed Network** node hierarchy table. The **Managed Network** node hierarchy table is filtered to list only the managed devices that are currently under upgrade.

Click **Show details** in the controller upgrade in progress pop-up to view the upgrade status of the managed devices and their access points. The upgrade status of the managed devices is displayed as either **Controller installation has started and is in progress** or **Controller installation completed successfully**.

The upgrade status of the access points connected to the managed devices is displayed as a horizontal bar chart and can be switched to table view. The table view lists the following columns:

Table 191: *Upgrade Status of Access Points Connected to Managed Device*

Parameter	Description
Name	Specifies the name of the access point.
IP ADDRESS	Specifies the IP address of the access point.
MAC ADDRESS	Specifies the MAC address of the access point.
AP GROUP	Specifies the AP group of the access point.
STATUS	Specifies the status of upgrade (pending, installation successful, or reboot in progress) of the access point.

Scheduling Upgrade of Clusters

The following procedure describes how to schedule the upgrade of multiple clusters.

1. In the **Managed Network** node hierarchy, navigate to **Maintenance > Software Management**.
2. Select one or multiple clusters from the table. The table only lists the clusters and not the cluster members. The table lists the following columns:

Table 192: *Schedule Upgrade of Multiple Clusters*

Parameter	Description
NAME	Specifies the name of the cluster. The number displayed next to the cluster name, as a link, indicates the total number of cluster members. Hovering on the link opens a pop-up which lists the names of the cluster members.
CURRENT VERSION	Specifies the ArubaOS release version (release number followed by build number) that is already installed on the cluster. If cluster members run different versions of ArubaOS, this column displays the ArubaOS version that is installed on the cluster leader as a link. Hovering on the link opens a pop-up, which lists of names of the cluster members and corresponding ArubaOS release version.

Parameter	Description
	If an upgrade is in progress, this column displays the **Installation in progress** link. Hovering on the link opens a pop-up, which lists: ■ Additional information of the upgrade through the **Show details** link. ■ The time (in most significant non-zero units) when the upgrade was initiated. If an upgrade was initiated within the last 7 days, this column displays an upgrade status icon next to the value of the current version. Hovering on the upgrade status icon opens a pop-up, which lists: ■ Additional information of the upgrade through the **Show details** link. ■ The time (in most significant non-zero units) when the upgrade succeeded or failed. If an upgrade is scheduled, this column displays an upgrade schedule clock instead of the upgrade status icon next to the value of the current version. Hovering on the upgrade schedule clock opens a pop-up, which lists: ■ The ArubaOS release version and scheduled date and time. ■ An option to reschedule the date and time. ■ An option to cancel the scheduled upgrade.
ACCESS POINTS	Specifies the total number of access points that are configured on the cluster.
GROUP	Specifies the parent group of the cluster.

3. In the **INSTALLATION SETTINGS > When** menu:
   - To upgrade the selected clusters immediately, select **Now**.
   - To schedule the upgrade of the selected clusters at a later date and time, select **Later**. Use the date and time picker to schedule the date and time. If the upgrade is scheduled to occur at a later date and time, a **Pending changes will automatically be deployed on managed controllers when the software installation will start** message is displayed.



- You can specify any date and time from the current date and time until the next 30 days.
- The upgrade occurs at the scheduled local time of the cluster leader and not at the local time of the Mobility Master.

4. In the **Specify image file location, name and protocol to use for transfer > Server IP address** menu, specify the hostname, IPv4, or IPv6 address of the file server that hosts the ArubaOS image.
5. In the **Specify image file location, name and protocol to use for transfer > Image path** menu, specify the path of the ArubaOS image on the file server that hosts the image. You may use the **.** wildcard character to specify the default path.
6. In the **Specify image file location, name and protocol to use for transfer > Protocol** menu, select the protocol to transfer the ArubaOS image from the file server to the selected clusters.

Table 193: *Protocol to Transfer ArubaOS Image*

Parameter	Description
FTP	Uses the FTP protocol to transfer the ArubaOS image from the file server to the selected clusters. If you choose the FTP protocol, specify the username and password to log in to the FTP server.
SCP	Uses the SCP protocol to transfer the ArubaOS image from the file server to the selected clusters. If you choose the SCP protocol, specify the username and password to log in to the SCP server.
TFTP	Uses the TFTP protocol to transfer the ArubaOS image from the file server to the selected clusters.

7. In the **Specify image file location, name and protocol to use for transfer > Software to install** menu, specify the software release number of the ArubaOS image in release number-build number format.
8. In the **Specify image file location, name and protocol to use for transfer > Partition** menu, select the partition of cluster members where the ArubaOS image should be installed.

Table 194: *Partition of Managed Device*

Parameter	Description
auto	Upgrades the default boot partition of the selected managed devices.
partition 0	Upgrades partition 0 of the selected managed devices.
partition 1	Upgrades partition 1 of the selected managed devices.

9. Click **Install**. A **Install Confirmation** pop-up displays the **Network performance will be temporary impacted during software installation Do you want to install software now?** message.
10. In the **Install Confirmation** pop-up, click **Install**.

To schedule the upgrade of a cluster using the CLI, see :

```
(host) [mynode] #lc-cluster <cluster_prof> schedule upgrade <version> <year> <month>
<day> <hh> <mm> <ss>upgrade managed-devices copy configured-fileserver version 8.4.0.0
all schedule 2018 12 25 12 30 00
```

To reschedule the scheduled upgrade of a cluster using the CLI:

```
(host) [mynode] #lc-cluster v4 re-schedule upgrade <version> <year> <month> <day> <hh>
```

To cancel the scheduled upgrade of a cluster using the CLI:

```
(host) [mynode] #lc-cluster <cluster_prof> abort scheduled-upgrade
```

Viewing Status of Scheduled Upgrade of Multiple Clusters

A scheduled upgrade icon in the banner indicates that one or more clusters are scheduled for upgrade. Hovering over the scheduled upgrade icon opens a pop-up which lists the number of clusters that are scheduled for upgrade. Click the number in the pop-up to open the **Managed Network** node hierarchy table. The **Managed Network** node hierarchy table is filtered to list only the clusters that are scheduled for upgrade.

If the upgrade of clusters is in progress, a controller upgrade icon replaces the scheduled upgrade icon. Hovering over the controller upgrade icon opens a pop-up, which lists the number of clusters that are currently under upgrade. Click the number in the pop-up to open the **Managed Network** node

hierarchy table. The **Managed Network** node hierarchy table is filtered to list only the clusters that are currently under upgrade.

Click **Show details** in the controller upgrade in progress pop-up to view the upgrade status of the cluster members and their access points or all access points in the cluster. The upgrade status of cluster members is displayed as a table. The table view lists the following columns:

Table 195: *Upgrade Status of Multiple Clusters*

Parameter	Description
Name	Specifies the name of the cluster member.
IP ADDRESS	Specifies the IP address of the cluster member.
MAC ADDRESS	Specifies the MAC address of the cluster member.
STATUS	Specifies the status of upgrade (pending, installation successful, or reboot in progress) of the cluster member.

When a cluster member is selected in the table, the upgrade status of the access points connected to that cluster member is displayed as a horizontal bar chart and can be switched to table view. The table view lists the following columns:

Table 196: *Upgrade Status of Access Points in Cluster*

Parameter	Description
Name	Specifies the name of the access point.
IP ADDRESS	Specifies the IP address of the access point.
MAC ADDRESS	Specifies the MAC address of the access point.
AP GROUP	Specifies the AP group of the access point.
TARGET CONTROLLER	Specifies the name of the cluster member to which the access point is connected.
STATUS	Specifies the status of upgrade (pending, installation successful, or reboot in progress) of the access point.

When a cluster member is not selected in the table, the upgrade status of all access points in the cluster is displayed as a horizontal bar chart and can be switched to table view.

Backing Up the Flash File System

You can back up the flash file system using the WebUI or CLI.

In the WebUI

To back up the flash file system:

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Configuration Management > Backup**.
2. Select **Flash** in **Select what to backup**.
3. Click **Create Backup**. By default, the flash file system is backed up to a flashbackup.tar.gz file.

In the CLI

To back up the flash file system, execute the following command:

```
(host) [mynode] #backup flash
```

Copying the Flash File System

Before copying the flash file system, back up of the flash file system. To back up the flash file system, perform the steps listed in Copying the Flash File System.

To copy the flash file system using the WebUI:

1. In the **Mobility Master** node hierarchy, navigate to **Diagnostics > Technical Support > Copy Files**.
2. Select **Flash file system** in **Select source file**.
3. Select **flashbackup.tar.gz** in **File name**.
4. Select the destination file in **Select destination file**.
5. Specify a file name in **File name**.
6. Click **Copy**.

Backing Up the Configuration

To back up the configuration using the WebUI:

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Configuration Management > Backup**.
2. Select **Configuration** in **Select what to backup**.
3. Click **Create Backup**. By default, the flash file system is backed up to a configbackup.tar.gz file.

Restoring the Flash File System

The flash file system can be restored it is previously backed up. To back up the flash file system, perform the steps listed in Backing Up the Flash File System.

To restore the flash file system using the WebUI:

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Configuration Management > Restore**.
2. Select **Flash** in **Select what to restore**.
3. Click **Restore**.

Restoring the Configuration

The configuration can be restored it is previously backed up. To back up the configuration, perform the steps listed in Backing Up the Configuration.

To restore the configuration using the WebUI:

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Configuration Management > Restore**.
2. Select **Configuration** in **Select what to restore**.
3. Click **Restore**.

Clearing the Configuration

To clear configuration using the WebUI:

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Configuration Management > Clear Configuration**.
2. Click **Clear Configuration**.

Synchronizing the Database

To synchronize the Database using the WebUI:

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Configuration Management > Synchronize Database**.
2. Click **Synchronize**.

Exporting the WMS Database

To export the WMS database using the WebUI:

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Configuration Management > WMS Database**.
2. Click **Export**.
3. Specify a file name in the **Export file** window.
4. Click **OK**.

Importing the WMS Database

To import the Database using the WebUI:

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Configuration Management > WMS Database**.
2. Click **Import**.
3. Specify a file name in the **Import file** window.
4. Click **OK**.

Clearing Old Entries in the WMS Database

To clear the old entries in the WMS database using the WebUI:

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Configuration Management > WMS Database**.
2. Click **Clean Old Entries**.
3. Click **OK** in the clean WMS DB window.

Re-initializing the WMS Database

To re-initialize the WMS database using the WebUI:

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Configuration Management > WMS Database**.
2. Click **Re-initialize Database**.
3. Click **OK** in the WMS database re-initialization window.

Configuring AP Image Preload

The AP image preload can be configured only on stand-alone controllers.

To configure the AP image preload using the WebUI, complete the following steps:

1. In the **Managed Network** node hierarchy, navigate to the **Maintenance > Access Point > Preload Image** tab.
2. Enable the **Ap image preload** toggle switch.
3. Select the partition from where you wish to download the image. For example, **Partition 0** or **Partition 1**.
4. Specify the maximum number of simultaneous download in **Maximum number of simultaneous download**.
5. To preload all APs, select **All APs** in **APs to preload**.
6. To preload a specific set of APs, select **Specific APs** in **APs to preload**.
   - To preload an AP group, click the **Add** button at the bottom of the **AP GROUP** table.
     1. Specify the AP group in **Type the AP group** or select the AP groups in the **Add AP Group** window.
     2. Click **OK**.
   - To preload an AP, click the **Add** button at the bottom of the **AP NAME** table.
     1. Specify the AP in **Type the AP name** or select the AP names in the **Add AP Name** window.
     2. Click **OK**.
7. Click **Apply**.

Rebooting an AP

An AP can be rebooted only on stand-alone controllers.

To reboot an AP using the WebUI:

1. In the **Mobility Master** node hierarchy, navigate to **Maintenance > Access Point > Reboot**.
2. Select the AP in the **Access Point** table.
3. Click **Reboot**.

Managing Certificates

The Mobility Master 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.

This section describes the following topics:

- About Digital Certificates on page 973
- Obtaining Server Certificate on page 973
- Obtaining Client Certificate on page 975
- Importing Certificates on page 975
- Viewing Certificate Information on page 976
- Imported Certificate Locations on page 976
- Checking CRLs on page 977
- Certificate Expiration Alert on page 977
- Support for Certificates on USB Flash Drives on page 978

Starting from ArubaOS 8.0.1.0, Mobility Master and managed devices generate a default certificate (controller-issued server certificate) to demonstrate the authentication of the managed device for captive portal and WebUI management access while booting. The controller-issued server certificate is used as the default certificate for WebUI authentication, 802.1X termination, and SSO.



The default-self-signed server certificate in ArubaOS 8.0.0.0 is changed to controller-issued server certificate in ArubaOS 8.0.1.0.

Aruba *strongly* recommends that you replace the default certificate with a custom certificate issued for your site or domain by a trusted CA. This section describes how to generate a CSR to submit to a CA and how to import the signed certificate received from the CA into the managed device.

The managed device supports client authentication using digital certificates for specific user-centric network services, such as AAA FastConnect, VPN (see Virtual Private Networks on page 459), and WebUI and SSH management access. Each service can employ different sets of client and server certificates.

During certificate-based authentication, the managed device provides its server certificate to the client for authentication. After validating the server certificate of the managed device, the client presents its own certificate to the managed device for authentication. To validate the client certificate, the managed device checks the CRL maintained by the CA that issued the client certificate. After validating the certificate of the client, the managed device can check the user name in the certificate with the configured authentication server (this action is optional and configurable).



To ensure that the clients are always connected to the captive portal page through SSL, you must create a bundle of chained certificates and concatenate the bundle to the signed server certificate as part of webserver configuration.

When using X.509 certificates for authentication, if a banner message has been configured on the managed device, it displays before the user can login. Click on the **Login** button after viewing the banner message to complete the login process.

About Digital Certificates

Clients and the servers to which they connect may hold authentication certificates that validate their identities. When a client connects to a server for the first time, or the first time since its previous certificate has expired or been revoked, the server requests that the client transmit its authentication certificate. The certificate of the client is then verified against the CA which issued it. Clients can also request and verify the authentication certificate of the server . 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 managed device checks the validity of client certificates using CRLs maintained by the CA that issued the certificate.

Digital certificates employ 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 the public key of party A.

Obtaining Server Certificate

Best practice is to replace the default server certificate in the managed device with a custom certificate issued for your site or domain by a trusted CA. To obtain a security certificate for the managed device from a CA:

1. Generate a CSR on the managed device.
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 certificate and public key of the CA.
4. Install the server certificate, as described in Obtaining Server Certificate on page 973.

There can be only one outstanding CSR at a time in the device. Once you generate a CSR, you need to import the CA-signed certificate into the device before you can generate another CSR.



If certificates should be obtained for multiple managed devices using the CSR method in a Mobility Master- Managed devices topology, then generate and import the CSR for each managed device separately. If CSR is generated for multiple managed devices at a time, the previously generated CSR will be overridden and the certificate import will fail.

The following procedure describes how to generate a CSR on the managed device.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Certificates** tab and expand the **CSR** accordion.
2. Enter the following information:

Table 197: *CSR Parameters*

Parameter	Description	Range
CSR Type	Type of the CSR. You can generate a certificate signing request either with an elliptic curve key, or with a RSA key.	EC or RSA
Curve name	Length of the private or public key for ECDSA. This is applicable only if **CSR Type** is `ec`.	secp256r1 or secp384r1
Key Length	Length of the private or public key for RSA. This is applicable only if **CSR Type** is `rsa`. **NOTE:** RSA-1024 is not permitted if the managed device is operating in the FIPS mode.	1024, 2048, or 4096
Common Name	Typically, this is the host and domain name, as in www.example.com.	—
Country	Two-letter ISO country code for the country in which your organization is located.	
State/Province	State, province, region, or territory in which your organization is located.	
City	City in which your organization is located.	
Organization	Name of your organization.	

Parameter	Description	Range
Unit	Optional field to distinguish a department or other unit within your organization.	
Email Address	Email address referenced in the CSR.	

3. Click **Generate New**.
4. Click **View Current** to display the generated CSR. Select and copy the CSR output between the BEGIN CERTIFICATE REQUEST and END CERTIFICATE REQUEST lines, paste it into an email and send it to the CA of your choice.

The following CLI commands generate a CSR.

1. Run the following command:

```
crypto pki csr {rsa key_len <key_val> |{ec curve-name <key_val>} common_name <common_
val> country <country_val> state_or_province <state> city <city_val> organization
<organization_val> unit <unit_val> email <email_val>
```



RSA-1024 is not permitted if the managed device is operating in the FIPS mode.

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 Client Certificate

You can use the CSR generated on the managed device 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 can import the following types of certificates into the managed device:

- 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 public key of client. (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



You cannot export certificates from the managed device.

The following procedure describes how to import certificates into the managed device:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Certificates** tab
2. Expand the **Import Certificates** accordion.
3. In the **Import Certificates** table click +.
4. For **Certificate Name**, enter a user-defined name.
5. For **Certificate Filename**, click **Browse** to navigate to the appropriate file on your computer.
6. If the certificate is encrypted, enter and repeat the passphrase.
7. Select the **Certificate Format** from the drop-down list.
8. Select the **Certificate Type** from the drop-down list.
9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

The following CLI command imports 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 managed device. Click **View** to display the contents of a certificate.

To view the contents of a certificate using the CLI, use the following commands:

Table 198: *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 managed device 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 managed device 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 managed device are displayed.

Imported Certificate Locations

Imported certificates and keys are stored in the following locations in flash on the managed device:

Table 199: *Imported Certificate Locations*

Location	Description
/flash/certmgr/trustedCAs	Trusted CA certificates, either for root or intermediate CAs. Best practices is to import the certificate for an intermediate CA, you also import the certificate for the signing CA.
/flash/certmgr/serverCerts	Server certificates. These certificates must contain both a public and private key (the public and private key must match). You can import certificates in PKCS12 and X509 PEM formats, but they are stored in X509 PEM DES encrypted format.
/flash/certmgr/CSR	Temporary CSRs that have been generated on the managed device and are awaiting a CA to sign them.
/flash/certmgr/publiccert	Public key of certificates. This allows a service on the managed device 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 managed device checks with the appropriate CA to make sure that the certificate has not been revoked.



The managed device does not support download of CRLs.

Certificate Expiration Alert

The certificate expiration alert sends alerts when installed certificates, which correspond to trust chains, OCSP responder certificates, and any other certificates installed on the device. By default, the system sends this alert 60 days before the expiration of the installed credentials. This alert is then repeated periodically on a weekly or biweekly basis. This alerts consist of two SNMP traps:

- wlsxCertExpiringSoon
- wlsxCertExpired

Chained Certificates on the Remote AP

Chained certificates on the Remote AP (that is, certificates from a multi-level PKI) need to be in a particular order inside the file. The certificate of a Remote AP must be first, followed by the certificate chain in order, and then followed by the private key for the certificate. For example, with a root CA, a single intermediate CA, and a root CA, the PEM or PKCS12 file must contain the following parts, in this order:

1. Remote AP Certificate
2. Intermediate CA
3. Root CA
4. Private key



If this order is not followed, certificate validation errors occur. This order also applies to server certificates.

Support for Certificates on USB Flash Drives

This release now supports storing Remote AP certificates in a USB device. This ensures that the Remote AP certificate is activated only when the USB with the corresponding certificate is connected to the Remote AP. If the USB is removed from the Remote AP, the Remote AP certificate is deactivated and when the USB is connected to the Remote AP it acts a storage device and not as a 3G or 4G Remote AP.

The Remote AP supports only PKCS12-encoded certificates that are present in the USB. This certificate contains all the information that is required for creating the tunnel including the private key, Remote AP certificate with the chain of certificates, and the trusted CA certificate. There is a limit of three supported intermediate CAs.

Ensure you adhere to the following file naming guidelines when you are saving the certificate:

- The first twelve characters of the certificate file name should be the MAC address of the Remote AP. For example, if eth0 MAC address of a Remote AP is 00:0b:86:c2:00:6c, then the file name will be 000B86C2006C.P12 or 000B86C2006C_rap155.p12
- All alphabets of the MAC address in the file name should be in upper case.
- The file name can have additional characters after the MAC address separated by "_" for the purpose of identification.

If this naming convention is not followed a error will occur during certificate validation.

Follow the steps below to configure the USB certificate store:

1. Copy the PKCS12 certificate bundle to a USB device.
2. Enter a name for the certificate using the correct naming convention as mentioned above.



In the USB connected to the Remote AP, delete any duplicate **<mac-address>.p12** certificate file. Only one such file must be present in the USB.

If you unplug the USB device the Remote AP will become unresponsive. Reboot the Remote AP to bring it up with a custom certificate, if the USB device was unplugged.

Marking the USB Device Connected as a Storage Device

If the AP provisioning parameter “usb-type” contains the value “storage,” this indicates that the Remote AP will retrieve certificates from the connected USB flash drive.

Remote AP Configuration Requirements

The Remote AP needs to have one additional provisioning parameter, the pkcs12_passphrase, which can be left untouched or can store an ACSII string. The string assigned to this parameter is used as the passphrase for decoding the private key stored.



If you have an activated Remote AP that is using USB storage for the certificate, and you remove the USB storage, the Remote AP drops the tunnel. This is by design. However, for the Remote AP to re-establish the tunnel it has to be power cycled. It does not matter if you reinsert the USB storage before or after the power cycle as long as you power cycle it.

When the Remote AP successfully extracts all the information including the CA certificate, the Remote AP certificate and the Remote AP private key using the passphrase from the provisioning parameter, it successfully establishes the tunnel.

Certificate Enrollment Using EST

EST supports automatic enrollment of certificates with the EST Server. The certificates can now be enrolled or re-enrolled automatically by configuring an EST profile on the Controller.

Certificate Enrollment with EST allows users to use their own PKI instead of the factory or self-signed certificates available on the Controller or the AP. This enables the user to have maximum visibility and control over the management of the PKI used and address any issues related to security by themselves in a scaled environment.

This section describes the following topics:

- Configuring EST on the Controller
- Logging and Debugging
- Deployment Scenarios
- Using EST Certificates for Site-to-site VPN (IKEv2)

Configuring EST on the Controller

You can configure multiple EST profiles on a Controller, with different parameters using the CLI but only one will be activated using a global non-profile command.

This section contains the following topics:

- Important Points to Remember on page 979
- Prerequisites on page 979
- Configuring an EST Profile on page 980

Important Points to Remember

- For smooth deployment, EST activation should be done first on the MM and then on the MDs.
- EST server configuration should be common across all the Controllers deployed in the enterprise.

Prerequisites

Before configuring EST, ensure you complete the following prerequisites:

1. Import the CA or signing authority of EST server's SSL certificate on the Controller. For more information on importing certificates, refer to Managing Certificates on page 972.
2. Ensure time synchronization between all the devices involved in EST enrollment. For more information on time synchronization, refer to Clock Synchronization on page 1011.
3. If EST profile contains an FQDN as the server host, ensure that the DNS Server and domain name are configured on the enrolling devices. For information on configuring a DNS Server and a DNS name, refer to DHCP Address Pools on page 260.
4. If the EST server port is different from the default Port 443, ensure the corporate firewall allows the configured port.
5. Ensure that the server-host configured as part of the EST profile matches the Common Name or SubjectAltName fields of the EST Server’s certificate which is used during SSL handshake.
6. For Remote AP deployments, if the IPSEC inner pool address range is not a routable network within the enterprise domain, it is recommended to configure the route source nat rule so that

traffic gets srcnat with the Controller's IP address to reach the EST server. The route srcnat rule should be only to the EST server as the destination host and respective port number used as part of EST profile parameters. For more information on configuring route source NAT, refer to Enabling Remote AP Advanced Configuration Options on page 859.

7. When ClearPass Policy Manager is used as the EST server, the default EST services are enabled with the SHA512 RSA signature which is unsupported on the AP. The RSA settings must be changed to either SHA256 or SHA384 in order to enroll EST on both the AP and the Controller successfully.

Enhancements to EST Profile

Starting from ArubaOS 8.6.0.0, the following EST enhancements can be configured by the user,

- Users can configure the username and password for authentication. These credentials are used during the enrollment process and the server will use these credentials for authenticating the clients.



The Username/password authentication and the challenge-password authentication methods are mutually exclusive. Only one of the authentication methods can be used. CLI and WebUI will throw an error when both of the authentication methods are configured at the same time.

- Users can configure the optional parameter, Organizational Unit Name (OU) in the EST profile. If this field is configured, OU is inserted in the CSR and subsequently becomes part of the enrolled EST certificate.
- Users can configure arbitrary labels for EST enrollment and re-enrollment to perform different EST operations. The arbitrary label will be used for CA cert operations. The arbitrary enrollment label and the arbitrary re-enrollment label will be used for CSR Attributes operations. These two labels are optional parameters and if not configured the default arbitrary label will be used for enrollment and re-enrollment of EST server.
- EST client will use the already enrolled certs during re-enrollment.
- Users can change the credentials in an already activated EST profile and use the latest credentials without de-activating and re-activating the EST profile. This enhancement will avoid unnecessary AP reboot while changing the credentials. Only the username, password and challenge-password fields are allowed to change. Any change to the other profile parameters is not allowed.

Configuring an EST Profile

The following procedure describes how to configure a new EST profile.

1. Before configuring an EST profile, you must import the trusted CA to the Controller.
   a. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Certificates** tab.
   b. Click **+** in the **Import Certificates** section.
   c. Enter the name of the trusted CA in the **Certificate name** text box.
   d. Enter the certificate filename of the trusted CA in the **Certificate filename** text box. Click the **Browse** button to enter the full pathname.
   e. Enter an **Optional passphrase** and re-type the passphrase.

f. Select a certificate format from the **Certificate format** drop-down list. You can import certificates of format PEM, DER, and PKCS7.

g. Select **TrustedCA** from the Certificate type drop-down list.

h. Click **Submit**. The certificate appears in the **Import Certificates** section.

2. To configure a new EST profile on the Controller using the WebUI.

a. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Profiles** tab.

b. In the **All Profiles** menu, expand **EST profile > EST**.

c. In the **EST Profile: New Profile section**, click **+** to create a new profile.

d. Enter a name for the EST profile in the **Profile name** text box.

e. Enter the hostname of the EST server in the **Server host** text box.

f. The default Server port is 443. You may choose to enter a different EST server port in the Server port text box.

g. You can optionally enter a password in the **Challenge password** text box.

h. If you chose to enter a challenge password, retype the password in the **Retype** text box.

i. Enter an arbitrary label in the **Arbitrary label** text box.

j. Enter the certificate name of the EST server (same as in Step 1c) in the **Server's CA cert name** textbox.

k. Enter the Organizational Unit Name in the **Organizational Unit Name** text box.

l. Enter an arbitrary enrolment label in the **Arbitrary enrolment label** text box.

m. Enter an arbitrary reenrollment label in the **Arbitrary reenrollment label** text box.

n. Enter the **Username** and **password** for EST authentication.

o. Click **Submit**. the EST profile appears under the **EST Profile > EST** section of the **All Profiles** menu.

3. To complete EST enrollment on the Controller, you must activate the EST profile.

a. In the **Mobility Master** node hierarchy, navigate to the **Configuration > System > Certificates** tab.

b. Expand the **Enrollment over Secure Transport** accordion.

c. Set the Enable certificate provisioning using EST protocol toggle switch to active.

d. Select the EST profile from the **EST server** drop-down list.

e. Click **Submit**.

The following CLI commands configure a new EST profile.

```
(host) [mynode] (config)# est profile <profile_name>
(host) [mynode] (EST Profile <profile_name>)# arbitrary-label <arbitrary-label>
(host) [mynode] (EST Profile <profile_name>)# arbitrary-label-enrollment <arbitrary
enrollment label >
(host) [mynode] (EST Profile <profile_name>)# arbitrary-label-reenrollment <arbitrary
reenrollment label >
(host) [mynode] (EST Profile <profile_name>)# challenge-password <password>
(host) [mynode] (EST Profile <profile_name>)# clone <source>
```

```
(host) [mynode] (EST Profile <profile_name>)#organizational-unit-name <organizational-
unit-name>
(host) [mynode] (EST Profile <profile_name>)# server-host <IPv4 address/hostname>
(host) [mynode] (EST Profile <profile_name>)# server-port <port_number>
(host) [mynode] (EST Profile <profile_name>)# trustanchor-name <name>
(host) [mynode] (EST Profile <profile_name>)# username <username>
(host) [mynode] (EST Profile <profile_name>)# password <password>
(host) [mynode] (EST Profile <profile_name>)# end
```

The following CLI command activates an existing EST profile.

```
(host) [mynode] (config)# est-activate <profile_name>
```

Logging and Debugging

Execute the following command to enable the EST related logs on the Controller:

```
(host) [mynode](config)# logging security subcat est level debugging
```

Execute the following command to display all the logs related to EST operation on the Controller:

```
(host) [mynode]# show log security all [include|est]
```

Execute the following command to view the EST profile details:

```
(host) [mynode]# show est profile
```

Execute the following command to view the EST status of the Controller:

```
(host) [mynode]# show est status
```

Execute the following command to view the EST status of all the devices on the switches including the master:

```
(host) [mynode] (config)# show est status all
```

Execute the following command to view the details of the EST Certificate:

```
(host) [mynode]# show crypto pki <Cert_type> <Cert_name>
```

Execute the following command to view the IPsec map configuration on the Controller:

```
(host) [mynode]# show crypto-local ipsec-map
```



If EST certificate is used for the ipsec-maps, an EST certificate string is present in the output of the ipsec-map information.

Execute the following command to view the APs with EST enrollment:

```
(host) [mynode]# show ap database
```

Execute the following command to view the debug logs for the AP:

```
(host) [mynode]# show ap ap-cert-mgr log {<ap-name>|<ip-addr>}
```

Execute the following command to view the contents of the control plane security whitelist:

```
(host) [mynode]# show whitelist-db cpsec
```

Execute the following command to view the details of the Remote AP whitelist database:

```
(host) [mynode]# show whitelist-db rap
```

Deployment Scenarios

After EST enrollment, the factory certificate or self-signed based IPsec tunnel applicable for the following deployment scenarios will be replaced by the EST enrolled certificate for IPsec tunnel establishment:

Table 200: *Deployments with EST Enrollment*

Controller or AP			Authentication Method		EST
Managed Device	AP	Mobility Master	Factory Certificate	Hybrid Certificate	—
7xxx	—	7xxx	✓	x	✓
7xxx	—	x86	x	✓	✓
7xxx	AP	—	✓	x	✓
x86	AP	—	✓	x	✓
—	—	L2 MM Redundancy (72xx)	✓	x	✓



Deployment with EST enrollment is not supported for Cluster, HA, L3 MM redundancy, IPv6, Instant AP, External WhitelistDB topologies in the ArubaOS 8.2.0.0 release.

Using EST Certificates for Site-to-site VPN (IKEv2)

The following CLI commands enable the EST certificates for a Site-to-Site crypto map.

```
(host) [mynode] (config)# crypto-local ipsec-map <ipsec-map_name> <ipsec-map-number>
(host) [mynode] (config-submode)# enrolled-cert-auth
```

Configuring SNMP

Managed devices support versions 1, 2c, and 3 of SNMP for reporting purposes only. In other words, SNMP cannot be used for setting values in a system in the current ArubaOS version.



Aruba-specific MIBs describe the objects that can be managed using SNMP.

The following section provides information on configuring SNMP parameters:

SNMP Parameters

You can configure the following SNMP parameters:

Table 201: *SNMP Parameters*

Field	Description
Host Name	Host name of the managed device.
System Contact	Name of the person who acts as the System Contact or administrator for the managed device.
System Location	String to describe the location of the managed device.

Field	Description
Read Community Strings	Community strings used to authenticate requests for SNMP versions before version 3. **NOTE:** This is needed only if using SNMP v2c and is not needed if using version 3.
Enable Trap Generation	Enables generation of SNMP traps to configured SNMP trap receivers. Refer to the list of traps in the "SNMP traps" section below for a list of traps that are generated by the managed device.
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 managed device. Configure the following for each host or trap receiver: ■ IP address ■ SNMP version: can be 1, 2c, or 3. ■ Type: Trap or Inform (SNMPv2c or SNMPv3 only) ■ Engine ID: (SNMPv3 only) ■ Security string ■ UDP port on which the trap receiver is listening for traps. The default is the UDP port number 162. This is optional, and will use the default port number if not modified by the user.
If you are using SNMPv3 to obtain values from the managed device, you can configure the following parameters:	
User name	A string representing the name of the user.
Authentication protocol	An indication of whether messages sent on behalf of this user can be authenticated, and if so, the type of authentication protocol used. This can take one of the two values: ■ MD5: HMAC-MD5-96 Digest Authentication Protocol ■ SHA: HMAC-SHA-96 Digest Authentication Protocol
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 or decrypted with DES, the (private) privacy key for use with the privacy protocol.

The following procedure describes how to configure basic SNMP parameters.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > SNMP** tab.
2. If the managed device will be sending SNMP traps, click + in the Trap Receivers section to add a trap receiver.
3. If you are using SNMPv3 to obtain values from the managed device, click + in the **Users for SNMPv3** section to add a new SNMPv3 user.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure basic SNMP parameters.

```
hostname name
syscontact name
syslocation string
snmp-server community string
snmp-server enable trap
snmp-server engine-id engine-id
snmp-server host ipaddr version {1|2c|3} string [udp-port number]
snmp-server trap source ipaddr
snmp-server user name [auth-prot {md5|sha} password priv-prot DES password
```



Earlier versions of ArubaOS supported SNMP on individual APs. This feature is not supported by this version of ArubaOS.

MIB Files

To access ArubaOS MIB files:

1. Log in to Aruba Support site.
2. Navigate to **Download Software** > **ArubaOS**.
3. Navigate to the desired release folder.
4. Download the MIB file corresponding to the release.
5. Uncompress the MIB file to a local directory.

Enabling Capacity Alerts

Use the capacity alert feature to set managed device capacity thresholds which, when exceeded, will trigger alerts. The managed device will send a *wlsxThresholdExceeded* SNMP trap and a syslog error message when the managed device has exceeded a set percentage of the total capacity for that resource. A *wlsxThresholdCleared* SNMP trap and error message will be triggered if the resource usage drops below the threshold once again.

The following procedure describes how to configure the capacity thresholds:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > More** tab.
2. Expand the **Capacity Threshold** accordion.
3. Modify the capacity percentages for any of the thresholds described in Table 202.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following table describes the thresholds that can be configured with this feature.

Table 202: *Capacity Alert Thresholds*

Threshold	Description
Datapath CPU	Set an alert threshold for datapath CPU capacity. The <percentage> parameter is the percentage of the total datapath CPU capacity that must be exceeded before the alert is sent. The default threshold for this parameter is 30%.

Threshold	Description
Controlpath CPU	Set an alert threshold for controlpath CPU capacity. The <percentage> parameter is the percentage of the total controlpath CPU capacity that must be exceeded before the alert is sent. The default threshold for this parameter is 45%.
Controlpath memory	Set an alert threshold for controlpath memory consumption. The <percentage> parameter is the percentage of the total memory capacity that must be exceeded before the alert is sent. The default threshold for this parameter is 85%.
Total tunnels	Set an alert threshold for the tunnel capacity of the managed device. The <percentage> parameter is the percentage of the total tunnel capacity of the managed device that must be exceeded before the alert is sent. The default threshold for this parameter is 80%
Total users	Set an alert threshold for the user capacity of the managed device. The <percentage> parameter is the percentage of the total resource capacity that must be exceeded before the alert is sent. The default threshold for this parameter is 80%.
Total APs	The maximum number of APs that can be connected to a managed device is determined by the model type and licenses installed on that managed device. Use this command to trigger an alert when the number of APs currently connected to the managed device exceeds a specific percentage of its total AP capacity. The default threshold for this parameter is 80%.
Total locals	Set an alert threshold for the capacity to support branch and local managed devices on the master managed device. A master managed device can support a combined total of 256 branch and local managed device. The <percentage> parameter is the percentage of the total master managed device capacity that must be exceeded before the alert is sent. The default threshold for this parameter is 80%.

The following CLI command configures the capacity thresholds:
threshold

Sample Configuration

The following command configures a new alert threshold for datapath memory consumption:
(host) [mynode] (config) #threshold datapath-cpu 90

If this threshold is exceeded then subsequently drops below the 90% threshold, the managed device would send the following two syslog error messages.
May 14 13:13:58 nanny[1393]: <399816> <ERRS> |nanny| Resource 'Control-Path Memory' has gone above 90% threshold, value : 93
May 14 13:16:58 nanny[1393]: <399816> <ERRS> |nanny| Resource 'Control-Path Memory' has come below 90% threshold, value : 87

Configuring Logging

This section outlines the steps required to configure logging on a managed device.

For each category or subcategory of message, you can set the logging level or severity level of the messages to be logged. Table 203 summarizes these categories:

Table 203: *Software Modules*

Category or 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
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

Category or Subcategory	Description
dot1x	802.1X messages
radius	RADIUS user messages

For each category or subcategory, you can configure a logging level. Table 204 describes the logging levels in order of severity, from most to least severe.

Table 204: *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 managed device can direct these logs.

The following procedure describes how to configure the IP address of a syslog server to which the managed device can direct these logs.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Logging > Syslog Servers** page.
2. To add a logging server, click + in the **Syslog Servers** section.
3. Enter the **IP address** and the **Port number**.
4. Add the logging server to the list of logging servers. Ensure that the syslog server is enabled and configured on this host. Click **Apply**.
5. To select the types of messages you want to log, select **Logging Levels** .
6. Select the category or subcategory to be logged.
7. To select the severity level for the category or subcategory, select the level from the Logging Level drop-down list.
8. Select the logging format **CEF** or **BSD-standard** from the **Format** drop-down list.



The ArcSight CEF is a log management standard that uses a standardized logging format so that data can easily be collected and aggregated for analysis by an enterprise management system.

9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the IP address of a syslog server to which the managed device can direct these logs.

```
logging <ipaddr>
logging level <level> <category> [subcat <subcategory>]
```

Syslog operates over UDP and is connectionless. Therefore, it is not possible for the managed device to recognize a failure of the syslog server or the network path to the syslog server. By establishing an IPsec tunnel between the managed device and the syslog server, (see Planning a VPN Configuration) it is possible to indirectly track the status of the syslog server link.

After a failure occurs, the network administrator has to manually re-synchronize log files by copying them from the managed device to the syslog server. Use the **tar logs** CLI command to create an archive of all local logs, then use the **copy** CLI command to copy this archive to an external server. Log space is limited on the managed device, and depending on how long the outage lasted some local logs may be overwritten.

Syslog Files

To generate syslog file:

1. Log in to CLI of Mobility Master.
2. Switch to config mode.
3. Configure the logging command. Example: logging <ipv4addr> facility local0. For additional information, see ArubaOS 8.4.0.0 Command-Line Interface Reference Guide.
4. Issue the show logging command. For additional information, see ArubaOS 8.4.0.0 Command-Line Interface Reference Guide.

Enabling Guest Provisioning

The Guest Provisioning feature lets you manage guests who need access to the wireless network of your company. 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.

This section describes the following topics:

- Configuring Guest Provisioning Page
- Configuring Guest Provisioning User
- Creating Guest Accounts
- Optional Configurations

Configuring 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.

This section describes how to design a Guest Provisioning page in the WebUI using all three tabs.



You can create and design the Guest Provisioning page using the WebUI only.

See the following topics for more information:

- Configuring Guest Fields on page 990
- Configuring Page Design on page 992
- Configuring Email Messages on page 992

Configuring Guest Fields

The following procedure describes how to configure guest fields.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > Guest Provisioning** page. The Guest Provisioning configuration page displays with the Guest Fields section on top. This section contains the following columns:
   - Internal Name—The unique identifier that is mapped to the label in the UI.
   - Label in UI—A customizable string that displays in both the main listing pane and details sheet on the Guest Provisioning page.
   - Display in Details—Fields with selected check boxes appear in the Show Details popup-window.

     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 check boxes appear as columns in the management user summary page.
2. Select the check box next to each field, described in Table 205, that you want to appear on the Guest Provisioning page. Optionally, you can customize the label that displays in the UI.
3. Click **Preview Current Settings** in the **Guest Access** section to view what the Guest Provisioning page looks like while you are designing it.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.



Best practices is to check the **Display in Listing** field for only the most essential fields, so that the Guest Provisioning user does not have to scroll the guest listing horizontally to see all the columns.

Table 205: *Guest Provisioning—Guest Field Descriptions*

Guest Field	Description
guest_category	A guest is the person who needs guest access to the wireless network of the company. 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 company of the guest.
guest_email	Email address of the guest.
guest_phone	Phone number of the guest.
comments	Optional comments about the account status, meeting schedule, and so on of the guest.
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.
guest_status	Status a guest user. By default, this option is enabled.
grantor	The username of the person of who created the guest account.
grantor_role	The authentication role of the grantor.
sponsor_category	A sponsor is the primary contact of the guest for the visit. This is the label in the Guest Provisioning page for the sponsor information.
sponsor_username	Work department of the sponsor.
sponsor_email	Email address of the sponsor.
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

Guest Field	Description
optional_field_3	optional_field_2 description
optional_field_4	optional_field_2 description

Configuring Page Design

The Page Design section lets you specify the company banner, heading, and text and background colors that appear on the Guest Provisioning page.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > Guest Provisioning** tab and expand the **Page Design** accordion.
2. Click **Browse** to search for the Banner image



Best practices is to use a logo or banner image that is 600 x 100 pixels (width x height). The WebUI does not apply the size restrictions when you upload an image file, but the image is resized to 600 x 100 pixels when it displays or is printed.

3. Enter the label for the guest listing (the one you used in the Guest Fields tab) in the **Banner text** field.
4. Enter the hex value for the color of the text in the **Text Color** field. The text in the header of the guest listing displays in this color.
5. Enter the hex value for the color of the background in the **Background color** field. This determines the color of the header of the guest listing.
6. Click **Preview Current Settings** to view what the Guest Provisioning page looks like while you are designing it.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

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. Refer the following topics for more information:
   - Configuring SMTP Server and Port on page 992
   - Configuring SMTP server and port on page 993
2. Create the email messages using the WebUI. Refer the following topic for more information:
   - Creating Email Messages on page 993

Configuring SMTP Server and Port

The following procedure describes how to configure SMTP server and port.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > More** tab.
2. Click the **SMTP** section in the bottom of the page.

3. Enter the IP address of the SMTP server to which the managed device sends the guest provisioning email in the **IP Address of SMTP server** field.
4. Enter the number of the port through which the guest provisioning email passes in the **Port** field.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

Configuring SMTP server and port

The following command creates a guest-access email and sends guest user email through SMTP server IP address 1.1.1.1 on port 25.

```
(host) [md] (config) #guest-access-email
(host) [md] (Guest-access Email) #smtp-port 25
(host) [md] (Guest-access Email) #smtp-server 1.1.1.1
```

Creating Email Messages

The following procedure describes how to create email messages.

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > Guest Provisioning** page and expand the **Guest Email** or **Sponsor Email** accordion
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** check box when you want an email message to be sent to the guest or sponsor alerting them that a guest account has just been created.



Regardless of whether you select this option, the person responsible for managing the **Guest Provisioning** page may choose to send this email message manually at any time.

Figure 111 shows a sample email message that is sent to the guest after the guest account is created.

Figure 111 *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 **Submit**.

Configuring Guest Provisioning User

The guest provisioning user has access to the Guest Provisioning Page 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.

This section describes how to configure a guest provisioning user using the WebUI. All three methods are described below:

- Username and Password Authentication Method on page 994
- Static Authentication Method on page 994
- Smart Card Authentication Method on page 995

The following procedure describes how to set username and password authetication method.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Admin** tab.
2. To enable local authentication, click the **Enable Local Authentication** toggle switch.
3. In the **Management User** section, click **+**.
4. In the **Username** field, enter the name of the user who you want to configure as a guest provisioning user.
5. Enter the **Password** field, enter the password of the user.
6. In the **Retype Password** field, re-enter the password of the user.
7. From the **Role** drop-down list, select **guest-provisioning**.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following procedure describes how to set static authentication method:



Before using this method, make sure that the correct CA certificate is uploaded to the managed device.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Admin** tab.
2. In the **Management Users** table, click +.
3. In the **New User** section, select **Show users with certificate authentication**.
4. In the **Management Users with Certificate Authentication** table, click **+**.

5. Make sure that **External server** is not selected for **Authentication server.**
6. In the **Username** field, enter the name of the user who you want to configure as a guest provisioning user.
7. In the **Role** field, select **guest-provisioning** from the drop-down list.
8. Select the CA certificate you want to use from the **Trusted CA Certificate Name** drop-down list.
9. Enter client certificate serial number in the **Client Certificate Serial No.** field.
10. Click **Submit**.

The following procedure describes how to set smart card authentication method.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Admin** tab.
2. Expand the **Admin Authentication Options** accordion.
3. Under **Admin Authentication Options**, select **guest-provisioning** from the **Default role** drop-down list.
4. Select the **Enable** check box.
5. Select the server group from the **Server Group** drop-down list.
6. Click **Submit**.
7. Under **WebUI Authentication**, select the **Client Certificate** check box.
8. Click **Submit**.
9. Expand the **Management Users** accordion, click **Show users with certificate authentication**.
10. In the **Management Users with Certificate Authentication** table, click **+**.
11. Enable the **WebUI Certificate** toggle switch and select **External server** from **Authentication server** drop-down list.
12. Select the trusted CA certificate from the **Trusted CA Certificated Name** drop-down list.
13. Click **Submit**.
14. Click **Pending Changes**.
15. In the **Pending Changes** window, select the required check boxes and click **Deploy changes**.

In the CLI

This example creates a user named Alex and assigns her the role of guest provisioning.

```
(host) [md] (config)# mgmt-user Alex guest-provisioning
```

This example uses the CA certificate **mycertificate** with the serial number 1234 to authenticate user Laura in the guest provisioning role.

```
(host) [md] (config)# mgmt-user webui-cacert mycertificate serial 1234 Laura guest-provisioning
```

This example shows that using previously configured certificate (1234), authentication and authorization are automatically configured using an authentication server.

```
(host) [md] (config) #web-server profile
(host) [md] (Web Server Configuration) #mgmt-auth username/password certificate
(host) [md] (Web Server Configuration) #!
(host) [md] (config) #mgmt-user webui-cacert <certificate_name>
```

```
(host) [md] (config) #aaa authentication mgmt
(host) [md] (config) #server-group "internal"
(host) [md] (config) #mgmt-user webui-cacert default
(host) [md] (config) #mgmt-user webui-cacert 1234
```

The following section provides information on configuring the guest account information window.

Customizing Guest Access Pass

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. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services > Guest Provisioning** tab and expand the **Guest Access** accordion.
2. Click **Browse** to insert a logo or other banner information on the window.

NOTE

Best practices is to use a logo or banner image that is 600 x 100 pixels (width x height). The WebUI does not apply the size restrictions when you upload an image file, but the image is resized to 600 x 100 pixels when it displays or is printed.

3. You can enter text for the Policy portion of the window.
4. Click **Submit** to save your changes. Click **Preview current settings** to preview the window.

Creating Guest Accounts

After the Guest Provisioning user is created, that person can log in to the managed device using the preconfigured username and password. The Guest Provisioning page displays. (See Figure 113.) This is a sample page as the fields may differ based on how the network administrator designed the page.

Figure 112 *Creating a Guest Account—Guest Provisioning Page*



NOTE

If you do not want multiple guest users to share the same guest account concurrently, navigate to the Captive Portal Authentication and select the “Allow only one active user session” option. If a guest user authenticates successfully but the managed device detects there is already a guest session with the same guest username, the second login is rejected.

This section describes the following topics:

- Guest Provisioning User Tasks on page 996
- Importing Multiple Guest Entries on page 998

Guest Provisioning User Tasks

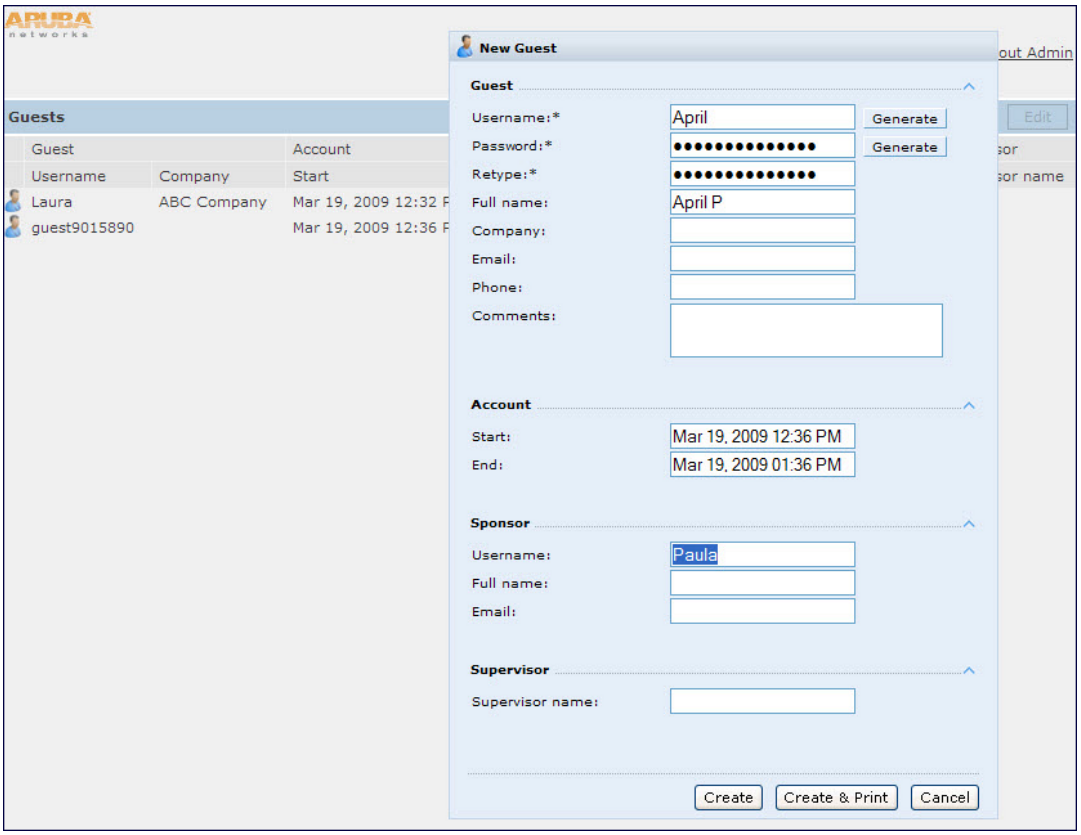
The Guest Provisioning user creates guest accounts by filling in information on the Guest Provisioning page. Tasks include creating, editing, manually sending email, enabling, printing, disabling and deleting guest accounts. The Guest Provisioning user can also manually send emails to either the guest or sponsor.

To create a new guest account, the Guest Provisioning user clicks **New** to display the New Guest window. (See Figure 113.) After filling in information into the fields, click **Create**. The guest account now displays on the Guest Provisioning page.

If you manually configure the user name and password, note the following:

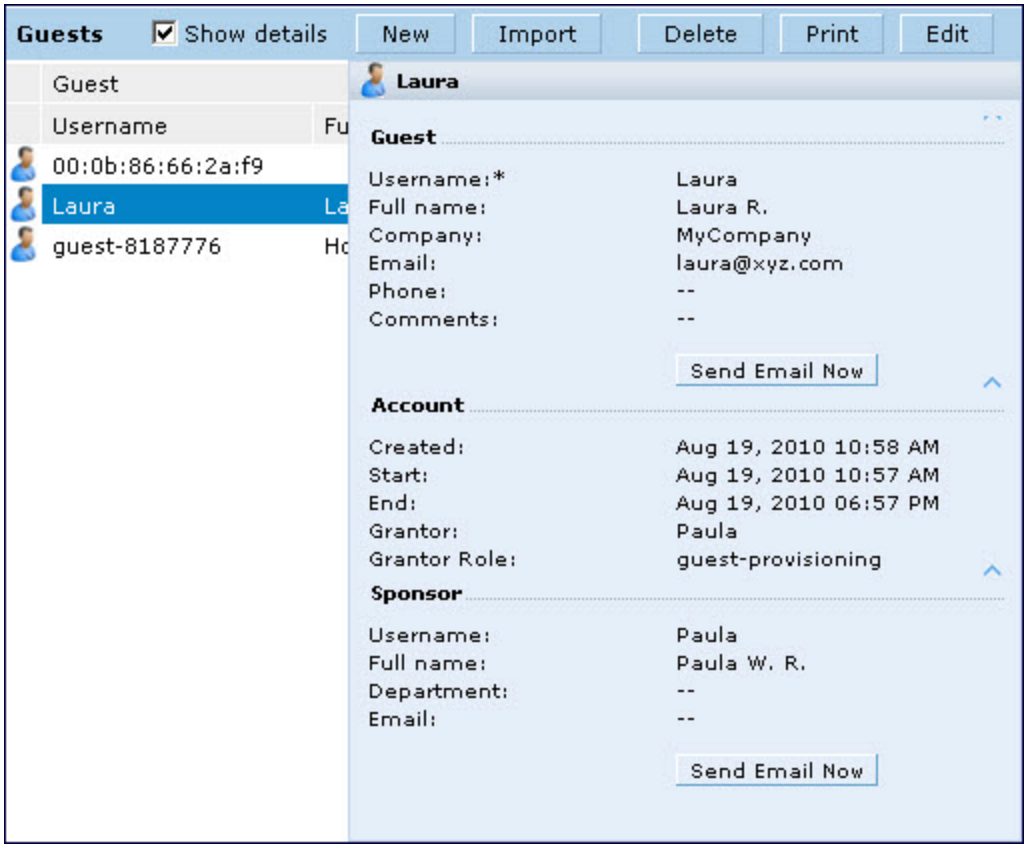
- User name entries support alphanumeric characters, however the percent sign (%) and trailing the back slash are not allowed.
- Passwords must have a minimum of six characters. You can use special characters for the password.
- Click on the **Account Start** and **End** fields to change the account start and end times. The default account start to end time setting is eight hours.

Figure 113 *Creating a Guest Account—New Guest Window*



To see details about an existing user account, highlight an existing account and select the **Show Details** check box. The Show Details popup-window displays. (See Figure 114.) The Guest Provisioning user can send out Email from this window to either the guest or the sponsor. When you send an email from the Details pop-up window, a pop-up message confirming that the email was successfully processed displays

Figure 114 *Creating a Guest Account—Show Details Pop-up Window*



Importing Multiple Guest Entries

The Guest Provisioning user can manually create individual guest entries, as previously described, or import multiple guest entries into the database from a CSV file. This is useful and more efficient if you want to enter multiple guest entries at once. To import multiple guest entries, you need to:

1. Create a CSV file that contains the guest entries
2. Import the CSV file into the database

Create a CSV file that contains multiple guest entries. Each field in an entry needs to be separated by a comma and each entry needs to end with a carriage return. The order of the fields is:

- First name of guest (required)
- Last name of guest (required)
- Email address of guest (optional)
- Phone number of guest (optional)
- User ID of guest (optional)
- Password of guest (optional)
- First name of sponsor (optional)
- Last name of sponsor (optional)
- Email address of sponsor (optional)

See Figure 115 for an example of how guest entries need to be formatted in a CSV file.

Figure 115 *CVS File Format—Guest Entries Information*

```
Gene,Phineas,gphineas@arubanetworks.com,(415)555-1212,guest-
gwang,abcdefg,Jane,Smith,jsmith@arubanetworks.com¶
Caulfield,Holden¶
John,Galt,,,guestlll0¶
```

Note the following limitations when creating guest entries in a CVS file:

- None of the field values can have a comma
- There is no format checking on field. Only the **local-userdb-guest** CLI command will validate proper format.
- Any extra columns, beyond the 9th column, are discarded.
- The WebUI only supports characters that the CLI supports.
- If a user ID of the guest is not provided, then it is automatically generated based on the numeric suffix in the Import Guest List window. See Figure 116.
- We recommend a maximum of 250 entries per CSV file.

To import a CSV file that contains multiple guest entries, the Guest Provisioning user must follow these steps:

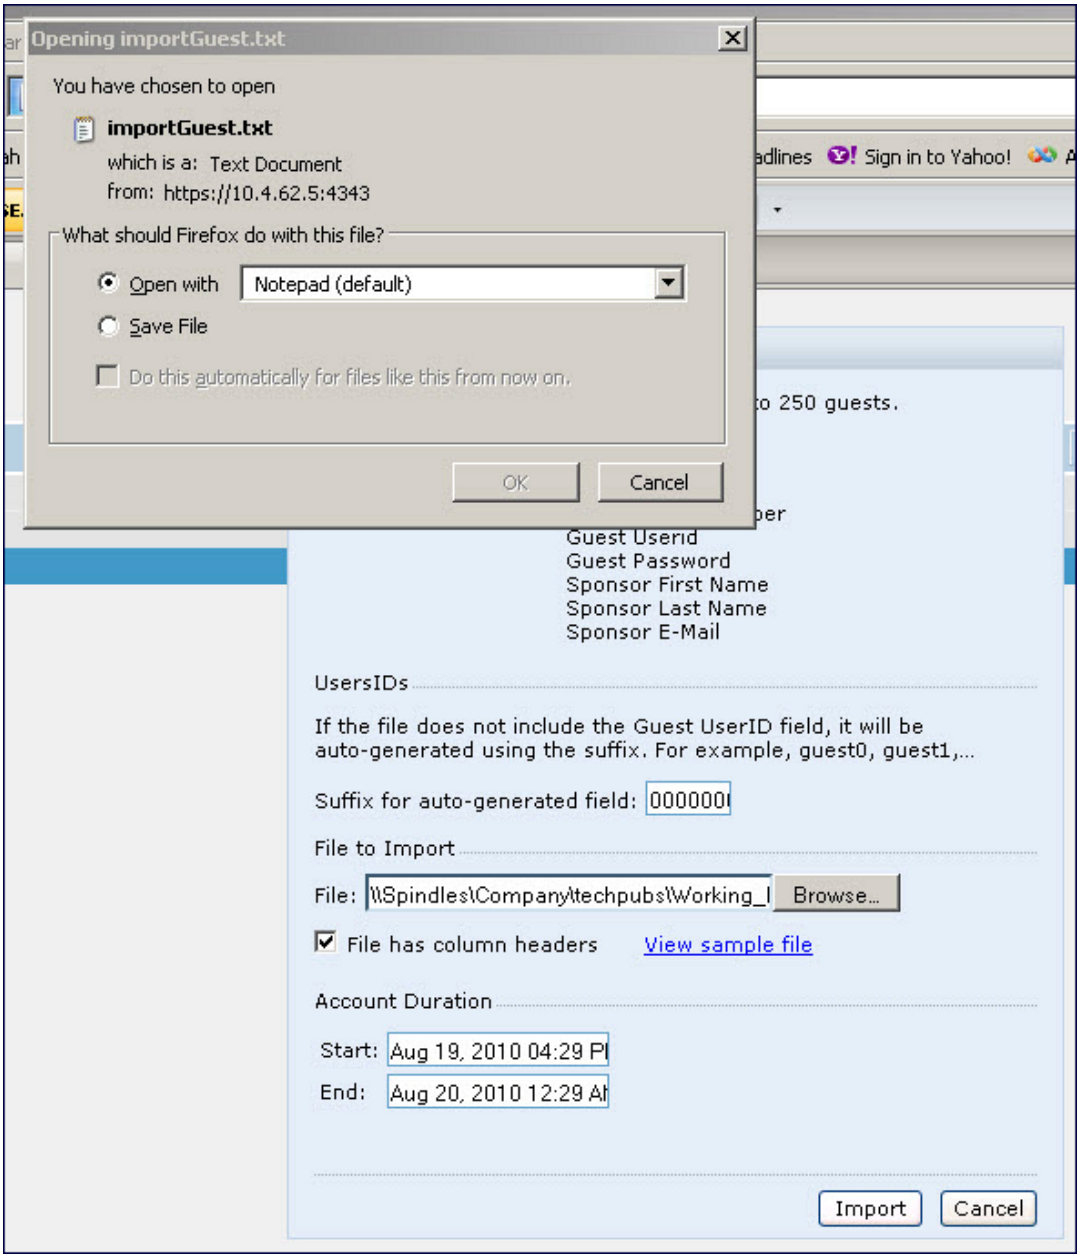
1. Log in to the WebUI using the username and password assigned to the Guest Provisioning user.
2. Click on **Import**. The **Import Guest List** pop-up window displays. See Figure 116.

Figure 116 *Importing a CSV file that contains Guest Entries*



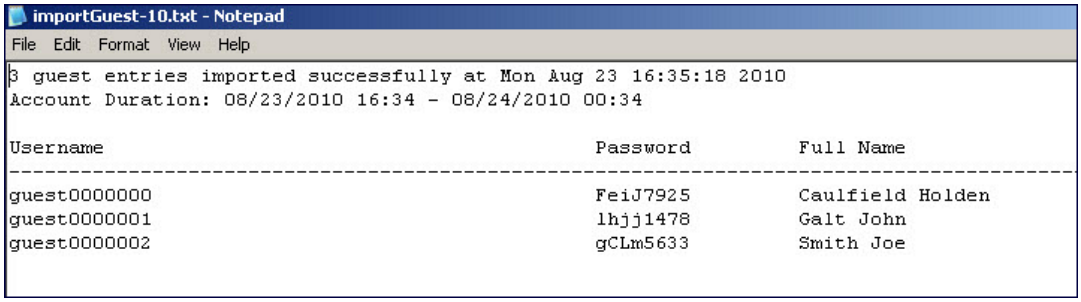
3. Click **Browse** to locate for the CSV file you want to import.
4. Click **Import**. A window displays that lets you open CSV file in text format. (See Figure 117.) Open the text file to see a summary of the number of users and error messages if users are not imported.

Figure 117 *Displaying the Guest Entries Log File*



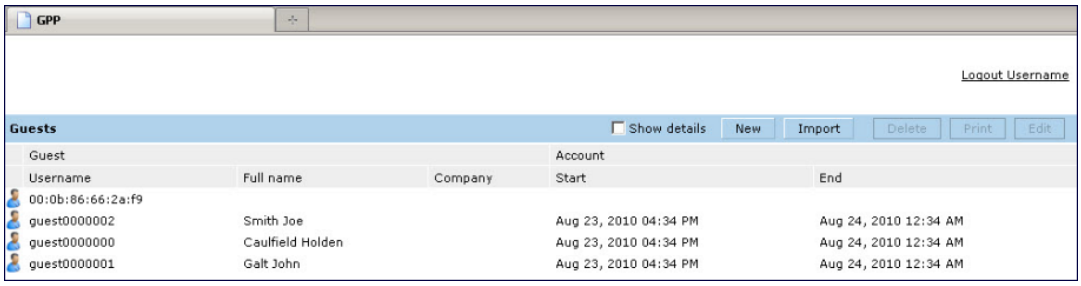
5. Open the text file. (See Figure 118.) Note that because no user ID is entered in the CSV file, a guest ID (username) is automatically generated based on the default value in the **Suffix for auto-generated** field. Make changes or corrections to the guest entry information in text file. A user can also change the start time and end time from this window. Save and exit the file.

Figure 118 *Viewing and Editing Guest Entries in the Log File*



6. Click **Cancel** to close the **Import Guest List** window. Guest entries are now displayed in the Guest Provisioning page.

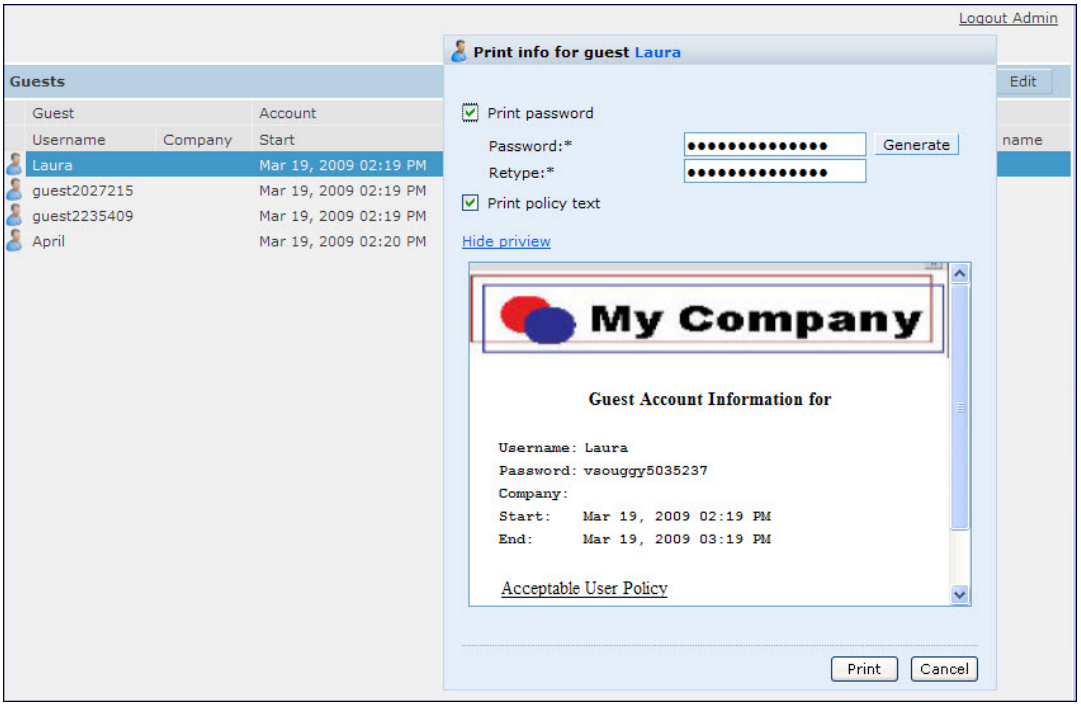
Figure 119 *Viewing Multiple Imported Guest Entries—Guest Provisioning Page*



To print guest account information:

1. Highlight the guest account you want to print and click **Print**. The **Print info for guest** window displays.
2. Click **Print password** if you want to print the guest password on the badge. Then enter or generate a new password for the guest. This modifies the existing guest password. (See Figure 120.)
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 120 *Printing Guest Account Information*



Optional Configurations

This section describes guest provisioning options that the administrator can configure. See the following topics for more information:

- Restricting one Captive Portal Session for each Guest on page 1003
- Setting Maximum Time for Guest Accounts on page 1004

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.

The following procedure describes how to restrict captive portal authentication for each guest:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Profiles** tab.
2. Select **Wireless LAN**.
3. Under **Wireless LAN**, select and open **Captive Portal Authentication**.
4. Add a new or select and existing profile.
5. Select the **Allow only one active user session** check box.
6. Click **Submit**.

7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command restricts captive portal authentication for each guest.

```
(host) [md] (config) #aaa authentication captive-portal <> single-session
```

Setting Maximum Time for Guest Accounts

You can set the maximum expiration time (in minutes) for guest accounts. If the guest-provisioning user attempt to add a guest account that expires beyond this time period, an error message is displayed and the guest account is created with the maximum time you configured.



If you set the maximum expiration time, it applies to all users in the internal database whether they are guests or not.

The following procedure describes how to configure maximum expiration time for guest accounts.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Security > Authentication** tab.
2. Select **Internal DB**.
3. Under Internal DB Maintenance, enter a value in **Maximum Expiration**.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures maximum expiration time for guest accounts.

```
(host) [md] (config) #local-userdb maximum-expiration <minutes>
```

Managing Files on Managed Device

- You can transfer the following types of files between the managed device and an external server or host:
- ArubaOS image file
- A specified file in the flash file system of the managed device or a compressed archive file that contains the entire content of the flash file system.



You back up the entire content of the flash file system to a compressed archive file, which you can then copy from the flash system to another destination.

- Configuration file, either the active running configuration or a startup configuration.
- Log files.
- Information about the Mobility Master or managed device for technical support.

You can use the following protocols to copy files to or from a managed device:

- FTP: Standard TCP or IP protocol for exchanging files between computers.
- TFTP: Software protocol that does not require user authentication and is simpler to implement and use than FTP.
- SCP: Protocol for secure transfer of files between computers that relies on the underlying SSH protocol to provide authentication and security.



You can use SCP only for transferring image files to or from the managed device, or transferring files between the flash file system on the managed device and a remote host. The SCP server or remote host must support SSH version 2 protocol.

The following table lists the parameters that you configure to copy files to or from a managed device.

Table 206: *File Transfer Configuration Parameters*

Server Type	Configuration
Trivial File Transfer Protocol (TFTP)	■ tftphost - tftp host IPv4 or IPv6 address ■ filename - absolute path of filename ■ flash: - copy to the flash file system ■ destination: - destination file name ■ system: - system partition ■ partition - partition 0 or partition 1
File Transfer Protocol (FTP)	■ ftphost - ftp server host name or IPv4 or IPv6 address ■ username - user name to log into server ■ filename - absolute path of filename ■ system: - system partition ■ partition - partition 0 or partition 1
Secure Copy (SCP) You must use the CLI to transfer files with SCP	■ scphost - scp host of IPv4 or IPv6 address ■ username - user name to secure to log into the server ■ filename - absolute path of filename (otherwise, SCP searches for the file relative to the home directory of the user) ■ flash: - copy to the flash file system ■ destfilename: - destination file name ■ system: - system partition ■ partition - partition 0 or partition 1

For example, you can copy an ArubaOS image file from an SCP server to a system partition on a managed device or copy the startup configuration on a managed device to a file on a TFTP server, You can also store the contents of the flash file system of the managed device 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 managed device or securely transfer a configuration file from flash to a remote host.

This section describes the following topics:

- Transferring ArubaOS Image Files on page 1005
- Backing Up and Restoring Flash File System on page 1006
- Copying Log Files on page 1007
- Creating Technical Support Report on page 1007
- Copying Other Files on page 1007

Transferring ArubaOS Image Files

You can download an ArubaOS image file onto a managed device 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 managed device, you must specify the system partition to which the file is copied. The WebUI shows the current content of the system partitions on the managed device. You have the option of rebooting the managed device with the transferred image file.

The following procedure describes how to transfer the ArubaOS image file.

1. In the **Mobility Master** node hierarchy, navigate to the **Maintenance > Software Management > Upgrade** page.
2. For **Upgrade using**, 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 managed device is to be rebooted after the image file is transferred, and whether the current configuration is saved before the managed device is rebooted.
6. Click **Upgrade**.

The following CLI commands transfer the ArubaOS image file.

```
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 Flash File System

You can store the entire content of the flash file system on a managed device 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.

Backing up Flash File System

The following procedure describes how to backup flash file system.

1. In the **Mobility Master** node hierarchy, navigate to the **Maintenance > Configuration Management > Backup** page.
2. Select **Flash**.
3. Click **Create Backup** to back up the contents of the flash system to the *flashbackup.tar.gz* file.
4. Click **Copy Backup** to enter the **Copy Files** page where you can select the destination server for the file.
5. Click **Copy**.

The following CLI commands backup flash file system.

```
backup flash
copy flash: flashbackup.tar.gz tftp: <tftphost> <destfilename>
copy flash: flashbackup.tar.gz scp: <scphost> <username> <destfilename>
```

Restoring Flash File System

The following procedure describes how to restore flash file system.

1. In the **Mobility Master** node hierarchy, navigate to the **Maintenance > Configuration Management> Restore** page.
2. Select **Flash**.
3. Click **Restore** to restore the *flashbackup.tar.gz* file to the flash file system.
4. Navigate to the **Maintenance > Software Management> Reboot** page.
5. Click **Reboot** to reboot.

The following CLI commands restore flash file system.

```
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.

The following procedure describes how to copy log files to an external TFTP or SCP server.

1. In the **Mobility Master** node hierarchy, navigate to the **Diagnostics > Technical Support > Copy Logs** page.
2. In the **Select destination file** field, select **TFTP server** or **FTP server** if you want to copy the log files in to a TFTP or FTP server. Enter the necessary destination details and the user credentials based on your server selection.
3. To download the log files into a WinZip file on your local PC, select **Download logs**.
4. Optionally, select **Include technical support info** if you want to include the support information in the logs.
5. Click **Copy**.

The following CLI commands copy log files to an external TFTP or SCP server.

```
tar logs
copy flash: logs.tar tftp: <tftphost> <destfilename>
copy flash: logs.tar scp: <scphost> <username> <destfilename>
```

Creating Technical Support Report

The following procedure describes how to create technical support report file.

1. In the **Mobility Master** node hierarchy, navigate to the **Diagnostics > Technical Support > System Information** page.
2. Enter a name for the file in the **Output file name** field and click **Create Report File**.
3. Once the technical support report is created, **Successfully Saved** message is displayed.
4. You can choose one the following options to get the technical support information:
   a. **Show Report**
   b. **Download Report**
   c. **Copy Report.**

For more more information refer Copying Other Files on page 1007.

To get technical support information on a managed device, repeat the steps above in the Managed Node hierarchy.

The following CLI commands gather a technical support report file:

```
show tech-support
<filename>
user
```

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 managed device is rebooted. It contains all options saved by clicking the **Submit** 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 or it can be copied to the local system.

- 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 or it can be copied to the local system.

The following procedure describes how to copy a file in the flash file system.

1. In the **Mobility Master** node hierarchy, navigate to the **Diagnostics > Technical Support > Copy Files** tab.
2. Select the source where the file or image exists.
3. The WebUI will automatically select **Flash File System** under the **Destination Selection** menu. You can also copy the files to **TFTP server, FTP server or Copy to local drive**.
4. Click **Copy**.

The following CLI commands copy a file in the flash file system.

```
copy startup-config flash: <filename>
copy startup-config tftp: <tftphost> <filename>
copy running-config flash: <filename>
copy running-config ftp: <ftphost> <user> <filename> [<remote-dir>]
copy running-config startup-config
copy running-config tftp: <tftphost> <filename>
```

SCP Server Support

SCP, which is based on SSH, is a tool to securely transfer files between hosts over the network. SCP uses SSH for authentication as well as data transfer. By using this functionality, clients can upload or download files from and to a server using SCP.

Typically, this functionality is of much use to customers that do not want to host a separate SCP server. Some customers need the ArubaOS controller, managed device, or Mobility Master to act as an SCP server as well, so that they can copy the files which are used in the general operation of the controller, managed device or Mobility Master— for example, files such as configuration, software upgrade images, and so on—from or to the controller, managed device, or Mobility Master flash. Additionally, for some customers, this functionality helps to easily manage their devices with their existing scripts or tools.

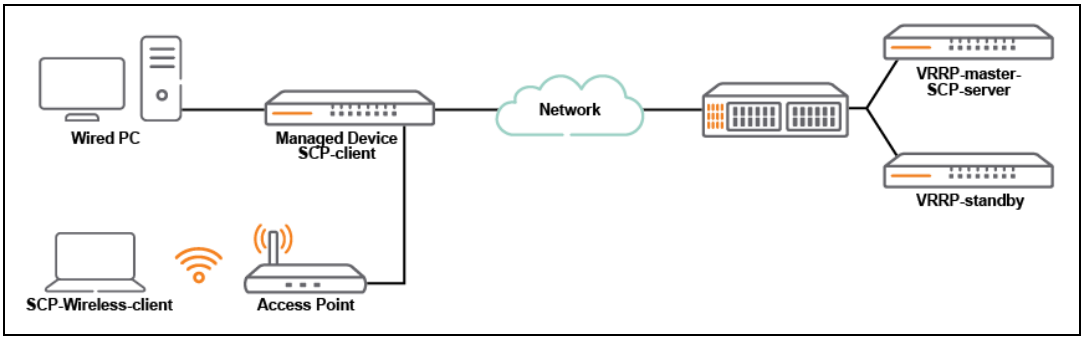
This section describes the following topics:

- Sample Topology on page 1008
- Configuring SCP Server Functionality on page 1009
- Verifying SCP Server Functionality Configuration on page 1009
- Debugging on page 1009
- Merits on page 1009
- Limitations on page 1010

Sample Topology

The following figure is a sample topology where VRRP-Master serves as the SCP server as well. A wireless client and a managed device in the network serve as the SCP clients. The Mobility Master must support enabling the SCP server functionality.

Figure 121 *Sample Topology - VRRP Master as SCP Server*



This is the sample topology for file transfer from external device to or from a controller that is serving as an SCP server.

NOTE

The maximum number of supported sessions is five at a time. So a customer's topology that has more than five managed devices cannot support the SCP server functionality.

Configuring SCP Server Functionality

To enable the SCP server functionality on the controller or managed device, execute the following command:

```
(host) [mynode] (config) #service scp
```

To disable the SCP server functionality on the controller or managed device, execute the following command:

```
(host) [mynode] (config) #no service scp
```

Verifying SCP Server Functionality Configuration

To view if the SCP server functionality on the controller or managed device is enabled or not, execute the following command:

```
(host) [mynode] #show scp
```

Debugging

To debug SCP functionality issues, execute the following command:

```
(host) [mynode]# show audit-trail
```

The output of the command shows all tracked file transfers, which can help in debugging any issue.

Merits

Enabling the SCP server functionality on a controller or managed device can help you perform the following tasks:

- Periodic backup of running configuration of the controller or managed device to another system.
- Update all the controllers or managed devices with a new configuration file.
- Upgrade all the controllers or managed devices in the network, without the need of an image server, by using the following steps:
    1. Copy the image to the flash storage of controller or managed device that serves as SCP server.
    2. In the controller or managed device, execute the **copy scp: <scp server ip> username system: partition 0|1** command.

NOTE

Aruba recommends to NOT use the controller or managed device that is serving as SCP server for a centralized image upgrade.

Limitations

Enabling the SCP server functionality in a controller or managed device can make it vulnerable to security issues. The following list briefs the restrictions enforced to use this SCP server functionality:

- Only SCP protocol is supported. SFTP or WinSCP protocols are not supported.
- You can copy only one file at a time. Directories are not supported.
- The SCP server functionality is available only in ArubaOS 8.x versions.
- The maximum number of simultaneous sessions supported is five.
- The implementation of this functionality is limited to work with APs that use the Beeliner interface.

Setting System Clock

You can set the clock on a managed device manually or by configuring the managed device to use a NTP server to synchronize its system clock with a central time source.

ArubaOS 8.2.0.0 introduces support for automatic timezone updates that include the relevant daylight savings time (DST) across timezones. This is done in view with keeping the time up-to-date and precise with daylight savings time adjustments effected automatically. Hence, the **Automatically adjust clock for Daylight Saving Time** check box that was available in earlier versions is no more available in the **Clock** accordion.

This section describes the following topics:

- Manually Setting Clock on page 1010
- Clock Synchronization on page 1011
- Configuring NTP Authentication on page 1011
- Timestamps in CLI Output on page 1012

Manually Setting Clock

The following procedure describes how to manually set the clock.

1. Login to the managed device.
2. Navigate to the **Configuration > System > General** page and expand the **Clock** accordion.
3. For **Time**, select **Manually** from the drop-down list.
4. To set the local time and date, click the **Modify Date and Time** button. The **Modify Date and Time** dialog box is displayed.
5. Enter the date and time in yyyy-mm-dd and hh:mm:ss formats, respectively, in the corresponding boxes. Note that hh is in 24-hr format. Click **OK**.
6. For **Time zone**, select the appropriate timezone from the list arranged in alphabetic order of countries.



An automatic update on the timezone picked and associated daylight savings time is available to keep the time up-to-date and relevant.

7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures the date and time.

```
clock set <year> <month> <date> <hour> <minutes> <seconds>
```

The following CLI command configures the time zone and daylight savings time adjustment.

```
clock timezone <WORD> <-23 - 23>
```



USE IANA time zone wording

Clock Synchronization

You can use NTP to synchronize the managed device to a central time source. Configure the managed device to set its system clock using NTP by configuring one or more NTP servers.

For each NTP server, you can optionally specify the NTP iburst mode for faster clock synchronization. The iburst mode sends up ten queries within the first minute to the NTP server. (When iburst mode is not enabled, only one query is sent within the first minute to the NTP server.) After the first minute, the iburst mode typically synchronizes the clock so that queries need to be sent at intervals of 64 seconds or more.



The iburst mode is a configurable option and not the default behavior for the managed device, 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.

You can use either WebUI or CLI to configure the system clock using NTP:

The following procedure describes how to configure the system clock using NTP.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > General > Clock** page.
2. For Time, select the **Using NTP** option.
3. Complete the steps for NTP servers and NTP authentication, if you require any change in these options.
4. For **Time zone**, select the appropriate timezone from the list arranged in alphabetic order of countries.



An automatic update on the timezone picked and associated daylight savings time is available to keep the time up-to-date and relevant.

5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures the system clock using NTP.

```
ntp server ipaddr [iburst]
```

Configuring NTP Authentication

The Network Time Protocol adds security to an NTP client by authenticating the server before synchronizing the local clock. NTP authentication works by using a symmetric key which is configured by the user. The secret key is shared by both the managed device and an external NTP server. This helps identify secure servers from fraudulent servers.

Starting from ArubaOS 8.2.1.0, a new NTP authentication option using SHA1 digest is available. A new parameter, **sha1**, is introduced in the **ntp authentication-key** command.

The following procedure describes how to configure NTP authentication.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > General > Clock** page.
2. To add a new NTP server, under **NTP Servers**, click **+**. Note that you can add up to 14 NTP servers.



You can drag and drop the required NTP servers in the **NTP Servers** table to reorder them in the list. Maintaining the correct order of the NTP servers determines which server will be contacted first.

3. In the **Add NTP Server** section, select **IPv4** or **IPv6** from the **IP version** drop-down list.
4. Enter the IPv4/IPv6 address of the NTP server in the **IPv4** or **IPv6** text box.
5. Select the **iburst mode** check box, if desired. It is disabled by default.
6. Enter the authentication key to be used by NTP server in the **Authentication key ID** text box. The range of allowed values is 1–65534; default is 1.
7. Click the **Use NTP authentication** toggle switch, if required. By default, the toggle switch is disabled.
8. Under **NTP Authentication Keys**, click **+**. The **Add NTP Authentication Key** section is displayed.
9. Enter the authentication key in the **Authentication key ID** text box. The allowed range of numeric values is 1–65534.
10. Select an hash algorithm from the **Hash algorithm** drop-down list. The two available options are md5 and sha1.
11. Enter the secret of the authentication key in the **Secret key** text box. The valid key value must be an ASCII string from 0 to 255 characters.
12. Select the **Trusted** check box to specify that the authentication key is trusted. By default, the check box is cleared.
13. Click **Submit**.
14. Click **Pending Changes**.
15. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands enable NTP authentication, add authentication secret keys into the database, specify a subset of keys which are trusted, and enable the iburst option.

```
(host) [md] (config) #ntp authenticate
(host) [md] (config) #ntp authentication-key <key-id> md5 <key-secret>
(host) [md] (config) #ntp trusted-key <key-id>
(host) [md] (config) #ntp server <ipaddr> <iburst> <key>
(host) [md] (config) #ntp server <server IP> <iburst key> <key>
```

The following CLI command configures the SHA1authentication key option.

```
(host)[mynode](config) #ntp authentication-key <keyid> sha1 <keyvalue>
```

The authentication key ID must be in the range of 1–65534. The key value must be up to 255 ASCII characters.

The **show ntp authentication-keys** command helps you verify the NTP authentication key type. The output of this command displays the SHA1 key type and the secret field (in encoded format), when SHA1 authentication is configured. The following example shows the output of the **show ntp authentication-keys** command:

```
(host) [mynode] # show  ntp authentication-keys
Key Id     Key Type     Secret
------     --------     ------
41          sha1        ********
```

Timestamps in CLI Output

The timestamp feature can include a timestamp in the output of each show command issued in the command-line interface, indicating the date and time the command was issued. Note that the output of **show clock** and **show log** do not include timestamps, even when this feature is enabled.

The following CLI command enables timestamp.

(host) [md] (config) #clock append

ClearPass Policy Manager Profiling with IF-MAP

This feature is used in conjunction with ClearPass Policy Manager. It sends HTTP User Agent Strings and mDNS broadcast information to ClearPass Policy Manager so that it can make more accurate decisions about what types of devices are connecting to the network.

The following procedure describes how to configure ClearPass Policy Manager.

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles > All Profiles > Other Profiles**.
2. Click the **CPPM IF-MAP** profile and click + in **CPPM IF-MAP** profile.
3. Configure this profile according to the following parameters:

Table 207: *ClearPass Policy Manager IF-Map Configuration Parameters*

Parameter	Description
CPPM IF-Map Interface	Enables the feature
Host IP address	IP address or hostname of the ClearPass Policy Manager IF-MAP server
Portnum	Enter a port number.
Username	Username for the user who performs actions on the ClearPass Policy Manager IF-MAP server. Range must be between 1-255 bytes in length.
Password	Password of the user who performs actions on the ClearPass Policy Manager IF-MAP server. Range between 6-100 bytes in length.

The following CLI commands configure ClearPass Policy Manager.

```
(host) [md] (config) #ifmap
(host) [md] (config) #ifmap cppm
(host) [md] (CPPM IF-MAP Profile) #server host <host>
(host) [md] (CPPM IF-MAP Profile) #port <port>
(host) [md] (CPPM IF-MAP Profile) #password <psswd>
(host) [md] (CPPM IF-MAP Profile) #enable
```

The following CLI command shows if the ClearPass Policy Manager interface is enabled and the status of the ClearPass Policy Manager server IP address, username, and password.

```
(host) [md] (CPPM IF-MAP Profile) #show ifmap cppm
CPPM IF-MAP Profile
-------------------
Parameter              Value
---------              -----
CPPM IF-MAP Interface  Enabled
CPPM IF-MAP Server     10.4.191.32:443 admin/********
```

The following CLI command shows the state of all enabled ClearPass Policy Manager servers.

```
(host) [md] (CPPM IF-MAP Profile) #show ifmap state cppm
CPPM IF-MAP Connection State [Interface: Enabled]
```

```
-----------------------------------------------
Server           State
------           -----
10.4.191.32:443  UP
```

Whitelist Synchronization

ArubaOS allows managed devices to synchronize their remote AP whitelists with the Aruba Activate cloud-based services. When you configure Activate whitelist synchronization, the managed device will securely contact the Activate server and download the contents of the whitelist on the Activate server to the whitelist on the managed device. The managed device and the Activate server must have layer-3 connectivity to communicate.

By default, this feature will both add new remote AP entries to the managed device whitelist and delete any obsolete entries on the managed device whitelist that were not on the Activate server whitelist. Select the add-only option to allow this feature to add or modify entries, but not delete any existing entries.

The following example enables the Activate whitelist service on the managed device. The **add-only** parameter allows only the addition of entries to the Activate remote AP whitelist database. This parameter is enabled by default. If this setting is disabled, the activate-whitelist-download command can both add and remove entries from the Activate database.

```
(host) [md] (config)# activate
(host) [md] (activate) #username
(host) [md] (activate) #password pass
(host) [md] (activate) #whitelist-enable
(host) [md] (activate) #add-only
```

The following command prompts the managed device to synchronize its remote AP whitelist with the associated whitelist on the Activate server:

```
(host) [md] (config) #activate whitelist download
```

Downloadable Regulatory Table

The downloadable regulatory table feature allows for the update of country domain options without upgrading the ArubaOS software version. A separate file, called the Regulatory-Cert, containing AP regulatory information will be released periodically on the customer support site. The Regulatory-Cert file can then be uploaded to a managed device and pushed to deployed APs.

The Regulatory-Cert includes the following information for each AP:

- All countries supported in the current release of ArubaOS (not just United States or Rest of World or any subset of countries)
- Allowed channels for each country
- Max EIRP for each channel and each country in the allowed list. The max values are specified for each PHY-type at which the AP is allowed to transmit on. The classified PHY-types are
  - 802.11 OFDM rates (802.11A or 802.11G mode)
  - 802.11b rates (CCK rates)
  - 802.11n HT20 and 802.11ac VHT20 rates (MCS0-7)
  - 802.11n HT20 and 802.11ac VHT40 rates (MCS0-7)
  - 802.11ac VHT80 rates
- DFS functionality for each channel and each country in the allowed list

This section describes the following topics:

- Important Points to Remember on page 1015
- Copying the Regulatory-Cert on page 1015
- Activating the Regulatory-Cert on page 1016
- Related Show Commands on page 1016

Important Points to Remember

- When a Regulatory-Cert is activated, the new file is checked against the default file built into ArubaOS. If the file is of a newer version, the activation is allowed. If the file is of a lower version, then the activation is not completed. The CLI of the managed device displays the following message upon failure:

```
(host) #ap regulatory activate reg-data-1.0_00002.txt
Failed to activate regulatory file reg-data-1.0_00002.txt. File Version should be
greater
than 1.0_43859
```

- APs do not rebootstrap or reboot on activation.
- If there is change in channel list or power level, APs will change the channel or power level. Impact is same as that of ARM channel or power change in that case.
- Clients are not disconnected upon regulatory file activation. Max latency impact during activation (with no channel changes) is less than 1s (applies to power change too).
- With channel change, the impact is similar to ARM channel change (depends on client behavior and if CSA is enabled or not).
- If support for the AP (Country) is added, the AP will move from AM to AP mode (if the AP is configured in AP mode of operation).

Copying the Regulatory-Cert

You can use the following protocols to copy the regulatory file to a managed device:

- FTP
- TFTP
- SCP

Additionally, regulatory files saved to a USB drive can be uploaded to a managed device equipped with a USB port.

The following procedure describes how to copy the Regulatory-Cert to the managed device.

1. In the **Mobility Master > host** node hierarchy, navigate to the **Diagnostics > Technical Support > Copy Files** tab.
2. Select the source (TFTP, FTP, SCP, or USB) where the file exists.
3. The managed device WebUI will automatically select **Flash File System** under the **Destination Selection** menu.
4. Click **Copy**.

The following CLI command downloads the regulatory file to the managed device:

```
copy
   ftp: <ftphost> <user> <filename>
   scp: <scphost> <username> <filename> flash: <destfilename>
   tftp: <tftphost> <filename> flash: <destfilename>
   usb: partition <partition-number> <filename> flash:  <destfilename>
```

The following CLI command shows the current regulatory and the content of the regulatory file.

```
show ap regulatory
```

```
show ap allowed-channels country-code <country-code> ap-type <ap-type
show ap allowed-max-eirp ap-name <ap-name> country-code <country-code>
show ap debug received-reg-table ap-name <ap-name>
```

Activating the Regulatory-Cert

Once the Regulatory-Cert has been added to the managed device, the new regulatory information must be activated and pushed to the APs.

The following CLI command activates a specific regulatory file loaded on the managed device.

```
ap regulatory activate <filename>
```

The following CLI command returns the regulatory file to the factory default regulatory-cert.

```
ap regulatory reset
```

In a Mobility Master-Managed Device deployment, the file syncing profile can be enabled to ensure that the regulatory-cert that is stored on the Mobility Master is shared with the managed devices.

File synchronization occurs when the downloadable regulatory table is activated in the Mobility Master and the Mobility Master synchronizes the downloadable regulatory table to all managed devices within the synchronization time and is automatically activated. File syncing is enabled by default, with a default sync time of 30 minutes. The sync time can be set between 30 to 180 minutes. To configure the file syncing profile, use the following commands

```
(host) [md] (config) #file syncing profile
(host) [md] (File syncing profile) #file-syncing-enable
(host) [md] (File syncing profile) #sync-time 30
```

Related Show Commands

The following CLI command shows the version of Regulatory Cert currently active on all managed devices.

```
(host) [md] #show switches regulatory
```

The following CLI command shows the file synching profile settings.

```
(host) [md] #show file syncing profile
```

Infrastructure for Supporting Database Upgrade

ArubaOS 8.2.0.0 introduces the infrastructure to support database upgrade. This feature exposes an API to the existing users of the database. The API takes in the name of the schema file associated with the application and upgrades its database.

This feature only supports the addition of new columns to the existing tables in the database. The upgrade of the schema will happen only during boot up and only during an image upgrade. Any subsequent reloads or restart of the application will not trigger the upgrade.



The downgrade of the database is not supported after the upgrade.

The API that is exposed to the applications when they want to upgrade their database:

```
int upgrade_postgres_db(char *sqlFile);
```

The database schema file name is passed to the API.

When a new image is downloaded for upgrade, and the system is not reloaded with that partition, then database upgrade is not triggered. The database for applications calling the database schema upgrade infrastructure API should be re-initialized in this case for the proper functioning of the application.

For example, the below steps will not trigger database upgrade:

1. Current system boot partition is 1.
2. Download new image to partition 0.
3. Change boot system to partition 1.
4. Reload.
5. Now change the boot system to partition 0.
6. Reload.

Configuring Concurrent Sessions

A check is added to limit the number of concurrent sessions that an administrator account can maintain. If the admin user tries to create a new session after the maximum concurrent user sessions limit is reached, then the system displays an error message and does not allow the user to login although the login credentials entered are valid.

The following CLI command restricts the maximum number of concurrent session for a management user:

```
(host) [mynode] (config) #mgmt-user <username> <rolename> max-concurrent-sessions
```

The following CLI command checks the maximum number of concurrent session for a management user:

```
(host) [mynode]#show management user
```

Implementing Management User Audits

When a user successfully logs in, the managed device must notify the administrator of the time, date, and the location of the user. Currently the information provided is not consistent, to overcome this issue the following user related information can be tracked in ArubaOS 8.4.0.0:

- Location of the last successful login indicates the SSH/WebUI IP address or console port.
- Date and time stamp of the last successful login.
- Number of unsuccessful attempts since the last successful login.
- Number of successful attempts over a period of time, which can be configured.

When the user logs in, the following pop-up is displayed in the **Dashboard** page:

- <username> (<role>)
- Last login: Mon Jul 16 15:21:50 2018 from 10.216.162.200
- There have been 3 failed login attempts since your last successful login
- 6 previous logins in the last 2 days

The following CLI command configures an audit period:

```
(host) [mynode] (config) #mgmt-user audit-period <audit-period>
```

The following CLI command shows the time, date, and location specific to a management user:

```
(host) [mynode] #show mgmt-user audit-info <username>
```

Implementing Password Validation

When a PSK based management user changes the password, a check is added to ensure that there is at least a difference of 8 characters between the new password and the old password. The new password

is encrypted, stored, and compared with the old password to increase the security and strength of the password.

The following procedure describes how to configure management password policy:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System >Profiles** tab.
2. In the **All Profiles** list, expand the **Other Profiles** menu, then select **Mgmt Password Policy**.
3. Enter a value between 0 to 32 in the **Minimum number of differing characters between passwords** field.
4. Click **Save**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy Changes**.

The following CLI command uses the old password when the password is being changed:

```
(host) [mynode] (config) #mgmt-user <username> <rolename> old-password
```

Maintaining Standard Mandatory Notice and Consent Banner

Starting from ArubaOS 8.4.0.0 the managed device must retain the Standard Mandatory Notice and Consent Banner on the screen until the administrator acknowledges the usage conditions and takes explicit actions to log on for further access.

The following procedure describes how to configure login banner text:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > System > Admin > Admin Authentication Options** accordion.
2. Select the **Banner has to be accepted** check box.
3. In the **Login banner text** field, enter details of what the user should view when they login.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy Changes**.
7. When the user logs in after the above steps are executed, a screen is displayed with the text entered in the **Login banner text** field.
8. Click **I Accept**.

The following CLI command forces the user to accept the details in the banner before logging in:

```
(host) [mynode] (config) #banner enforce-accept
```

Zeroizing TPM Keys

Zeroizing a cryptographic module involves erasing sensitive parameters such as electronically stored data, cryptographic keys, and critical security parameters from a controller or an AP to prevent disclosure of information if the equipment is permanently and irreversibly decommissioned.

The following CLI command erases the TPM content and renders the controller permanently inoperable.



Do not use this command prior to RMA, as it permanently decommissions a controller or an AP and voids any support or warranty entitlement.

```
(host) [mynode] (config) #zeroize-tpm-keys
The effect of the action you are about to execute is not reversible. Are you sure you
want to implement this function? Press 'y' to proceed : [y/n]: y
```

```
This action will void the warranty on the controller and nullify the RMA. Are you still
sure you want to do this?(y/n): y
You are about to wipe the contents of the TPM and render the controller permanently
inoperable. Are you ready to go ahead?(y/n): y
TPM keys have been zeroized. Please reload the controller.
```

The following CLI command erases the TPM content and renders the AP permanently inoperable.

```
(host) [mynode] (config) #ap zeroize-tpm-keys <ap-name>/<ip-address>/<ipv6-address>
You are about to execute a command which will make the AP inoperable and void the RMA.
Are you sure you want to proceed? [y/n]: y
TPM keys have been zeroized. Please reboot the AP.
```

The following CLI commands show the TPM initialization errors.

```
(host) [mynode] (config) #zeroize-tpm-keys
(host) [mynode] #show tpm errorlog
```

The following CLI command checks the TPM certificate installed on the controller.

```
(host) [mynode] #show tpm cert-info
```

Chapter 39
Hotspot 2.0

ArubaOS incorporates Passpoint technology from the WFA HotSpot 2.0 (Release 1) Technical Specification Version 1.0.0 to simplify and automate access to public Wi-Fi networks. Follow the procedures in this chapter to help mobile devices identify which access points in your hotspot network are suitable for their needs, and authenticate to a remote service provider using suitable credentials.



Throughout this document, all references to Hotspot 2.0 refer to HotSpot 2.0 (Release 1) Technical Specification Version 1.0.0. Recent versions of this technical specifications exist, but they are not supported in this version of ArubaOS.

Starting from ArubaOS 8.7.0.0, 530 Series, AP-555, and 570 Series access points support Hotspot 2.0.

Hotspot 2.0 is a WFA Technical Specification for the Wi-Fi Alliance Wi-Fi CERTIFIED PassPoint program based upon IEEE P802.11u-2011 that provides wireless clients with a streamlined mechanism to discover and authenticate to suitable networks, and allows mobile users the ability to roam between partner networks without additional authentication. ArubaOS supports Hotspot 2.0 with enhanced network discovery and selection. Clients can receive general information about the network identity, venue, and type via management frames from the AP. Clients can also query APs for information about the available IP address type (IPv4 or IPv6) of the network, roaming partners, and supported authentication methods, and receive that information in ANQP Information Elements from the AP.

Hotspot Profile Configuration Tasks

The following sections describe the procedure to configure profiles for hotspot feature.

- Configuring Hotspot 2.0 Profiles on page 1022
- Configuring Hotspot Advertisement Profiles on page 1026
- Configuring ANQP Venue Name Profiles on page 1028
- Configuring ANQP Network Authentication Profiles on page 1030
- Configuring ANQP Domain Name Profiles on page 1031
- Configuring ANQP IP Address Availability Profiles on page 1031
- Configuring ANQP NAI Realm Profiles on page 1032
- Configuring ANQP Roaming Consortium Profiles on page 1034
- Configuring ANQP 3GPP Cellular Network Profiles on page 1035
- Configuring H2QP Connection Capability Profiles on page 1036
- Configuring H2QP Operator Friendly Name Profiles on page 1037
- Configuring H2QP Operating Class Indication Profiles on page 1038
- Configuring H2QP WAN Metrics Profiles on page 1038

Access Network Query Protocol

ANQP is an Advertisement Protocol implemented using the GAS frames allowing any STA to query another STA about ANQP elements even before the association event.

ANQP Information Element

ANQP Information Elements are additional data that can be sent from the AP to any STA (including other APs) to provide identify of the APs network and service provider. The STA can query for which ANQP elements are available for being query using the Query List element within a query, in which case the AP will reply with the Capability List elements indicating what other ANQP elements can be queried. Here are the ANQP elements that can be returned in the capabilities List element

- Venue Name: the this information element defines the venue group and venue type.
- Domain Name: this information element specifies the APs domain name.
- Network Authentication Type: if the network has ASRA, this information element defines the authentication type being used by the hotspot network.
- Roaming Consortium List: this information element contains information identifying the network and service provider, whose security credentials can be used to authenticate with the AP transmitting this element.
- IP address Availability: this information element provides clients with information about the availability of IP address versions and types which could be allocated to those clients after they associate to the hotspot AP.
- NAI Realm: this information element identifies and describes a NAI realm accessible using the AP and the method that this NAI realm uses for authentication.
- 3GPP Cellular Network Data: this information element defines information for a 3GPP Cellular Network for hotspots that have roaming relationships with cellular operators.
- Connection Capability: this information element defines hotspot protocol and port capabilities to be sent in an ANQP information element.
- Operating Class: this information element defines the channels on which the hotspot is capable of operating.
- Operator Friendly Name: this information element allows the definition of a free-form text field that can identify the operator and additional information about the location.
- WAN Metrics: this information element provides hotspot clients information about access network characteristics such as link status, capacity and speed of the WAN link to the Internet.

Hotspot Profile Types

ArubaOS supports several different Hotspot 2.0 configuration profile types for defining ANQP information elements. The term H2QP is used to define profile that define Hotspot 2.0 specific Information Elements.

Table 208: *ANQP and H2QP Profiles referenced by an Advertisement Profile*

Profile	Description
Hotspot Advertisement profile	An advertisement profile defines a collection of ANQP and H2QP profiles. Each hotspot 2.0 profile is associated with one advertisement profile, which in turn references one of each type of ANQP and H2QP profile. For more information on configuring this profile, refer to Configuring Hotspot Advertisement Profiles on page 1026
ANQP 3GPP Cellular Network profile	Use this profile to define priority information for a 3GPP Cellular Network used by hotspots that have roaming relationships with cellular operators.

Profile	Description
	For more information on configuring this profile, refer to Configuring ANQP 3GPP Cellular Network Profiles on page 1035
ANQP Domain Name profile	Use this profile to specify the hotspot operator domain name. For more information on configuring this profile, refer to Configuring Hotspot Advertisement Profiles on page 1026
ANQP IP Address Availability profile	Use this profile to specify the types of IPv4 and IPv6 IP addresses available in the hotspot network. For more information on configuring this profile, refer to Configuring ANQP IP Address Availability Profiles on page 1031
ANQP NAI Realm profile	An APs NAI Realm profile identifies and describes a NAI realm accessible using the AP, and the method that this NAI realm uses for authentication. For more information on configuring this profile, refer to Configuring ANQP NAI Realm Profiles on page 1032
ANQP Network Authentication profile	Use the ANQP Network Authentication profile to define the authentication type used by the hotspot network. For more information on configuring this profile, refer to Configuring ANQP Network Authentication Profiles on page 1030.
ANQP Roaming Consortium profile	Name of the ANQP Roaming Consortium profile to be associated with this WLAN advertisement profile. For more information on configuring this profile, refer to Configuring ANQP Roaming Consortium Profiles on page 1034
ANQP Venue Name profile	Use this profile to specify the venue group and venue type information be sent in an ANQP information element in a GAS query response. For more information on configuring this profile, refer to Configuring ANQP Venue Name Profiles on page 1028.
H2QP Connection Capability profile	Use this profile to specify hotspot protocol and port capabilities. For more information on configuring this profile, refer to Configuring H2QP Connection Capability Profiles on page 1036
H2QP Operating Class Indication profile	Use this profile to specify the channels on which the hotspot is capable of operating For more information on configuring this profile, refer to Configuring H2QP Operating Class Indication Profiles on page 1038
H2QP Operator Friendly Name profile	Use this profile to define the operator-friendly name sent by devices using this profile. For more information on configuring this profile, refer to Configuring H2QP Operator Friendly Name Profiles on page 1037
H2QP WAN Metrics profile	Use this profile to specify the WAN status and link metrics for your hotspot. For more information on configuring this profile, refer to Configuring H2QP WAN Metrics Profiles on page 1038

Configuring Hotspot 2.0 Profiles

Use this profile to enable the hotspot 2.0 feature and define venue and OI settings for roaming partners. Each hotspot 2.0 profile also references an advertisement profile, which defines a set of ANQP or H2QP profiles which define other values for the hotspot feature. By default, hotspot 2.0 profiles reference the

default advertisement profile. For information on associating a different advertisement profile with a hotspot 2.0 profile, see Configuring Hotspot Advertisement Profiles on page 1026.

The following procedure describes how to configure a hotspot 2.0 profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the **All Profiles** table, expand **Wireless LAN**.
3. Expand **Hotspot 2.0**. The list of available hotspot 2.0 profiles appears in the **All Profiles** table.
4. Select an existing Hotspot 2.0 profile from the list of profiles or create a new profile by clicking **+**.
5. Select **Advertise Hotspot 2.0 Capability**.
6. Configure the following parameters as desired, then click **Submit** to save your settings.

Table 209: *Hotspot 2.0 Profile Settings*

Parameter	Description
Advertise Hotspot 2.0 Capability	This check box enables or disables the Hotspot 2.0 capability. When this feature is enabled, the IEs for this hotspot are included in beacons and probe responses from the AP and the ANQP Queries are answered to. This setting is disabled by default.
Enable Hotspot 2.0 OSEN	This setting enables OSEN Hotspot 2.0 to advertise and select an OSEN capable network. This setting is disabled by default.
ANQP Domain ID	Enter the ANQP Domain ID.
OSU NAI	Enter the OSU NAI (Network Access Identifier) for OSEN.
Use GAS Comeback Request/Response	By default, ANQP Information is obtained directly from a GAS Initial Response frame when size allows for it. If this parameter is enabled, advertisement information will not be included in the GAS Initial Response and will always force the usage of Comeback-Request and Comeback-Response frames. This option is disabled by default.
Additional Steps required for Access Enabled	This check box enables or disables the advertisement of ASRA in the Inter-networking information element. **NOTE:** If this parameter is enabled, the advertisement profile for this hotspot must reference a network authentication type profile.
Network Internet Access	This check box enables or disables the advertisement of Internet access in the Inter-networking information element.
Length of Query Response	The maximum number of 256 bytes that can be used for a GAS Initial-Response or GAS Comeback-Response frames. If the data exceeds this number, multiple Comeback-Response fragment will be used. The supported range is 1-256
Access Network Type	Specify the 802.11u network type. The default setting is public-chargeable. ■ **emergency-services**: emergency services only network ■ **personal-device**: personal device network ■ **private**: private network ■ **private-guest**: private network with guest access ■ **public-chargeable**: public chargeable network ■ **public-free**: free public network ■ **test**: test network ■ **wildcard**: wildcard network

Parameter	Description
Roaming Consortium OI value 1	Roaming consortium OI assigned to one of the top three roaming partners of the service provider. This additional OI will only be sent to a client if the OI parameter of the Additional Roaming Consortium is set to 1 or higher. **NOTE:** The service provider's own roaming consortium OI is configured using the ANQP Roaming Consortium profile.
Roaming Consortium OI value 2	Roaming consortium OI assigned to one of the top three roaming partners of the service provider. This additional OI will only be sent to a client if the OI parameter of the Additional Roaming Consortium is set to 2 or higher. **NOTE:** The service provider's own roaming consortium OI is configured using the ANQP Roaming Consortium profile.
Roaming Consortium OI value 3	Roaming consortium OI assigned to one of the top three roaming partners of the service provider. This additional OI will only be sent to a client if the OI parameter of the Additional Roaming Consortium is set to 3. **NOTE:** The service provider's own roaming consortium OI is configured using the ANQP Roaming Consortium profile.
Additional Roaming Consortium OI's (displayed in Advertisement Profile)	The number of additional ANQP Roaming Consortium profiles referenced by the Advertisement profile associated with this profile. The number must not include the OI defined within this Hotspot 2.0 profile.
HESSID	This optional parameter devices an AP's homogenous ESSID, which is that device's MAC address in colon-separated hexadecimal format.
Venue Group Type	Specify one of the following venue groups to be advertised in the IEs from APs associated with this hotspot profile. The default setting is unspecified. ■ assembly ■ business ■ educational ■ factory-or-industrial ■ institutional ■ mercantile ■ outdoor ■ reserved ■ residential ■ storage ■ unspecified ■ Utility-Misc ■ Vehicular **NOTE:** This parameter only defines the venue group advertised in the IEs from hotspot APs.
Venue Type	Specify a venue type to be advertised in the IEs from APs associated with this hotspot profile. The complete list of supported venue types is described in Venue Types on page 1029. This parameter only defines the venue type advertised in the IEs from hotspot APs.

Parameter	Description
PAME BI	This option enables the Pre-Association Message Exchange BSSID Independent bit, which is used by an AP to indicate whether the AP indicates that the Advertisement Server can return a query response that is independent of the BSSID used for the GAS Frame exchange.
Downstream Group Frames Forwarding Blocked	This option configures the Downstream Group Addressed Forwarding Disabled Mode. If this feature is enabled, it ensures that the AP does not forward downstream group-addressed frames. It is disabled by default, allowing the AP to forward downstream group-addressed frames.
Time Zone Format	The time zone in which the AP is operating, in the format `<std><offset>[dst[offset][,start[/time],end[/time]]` Where the <std> string specifies the abbreviation of the time zone, <dst> is the abbreviation of the timezone in daylight savings time, and the <offset> string specifies the time value you must add to the local time to arrive at UTC. **NOTE:** For complete details on configuring the timezone format, refer to section 8.3 of IEEE Std 1003.1, 2004 Edition.
Time Advertisement Capability	This parameter specifies the APs source of external time, and the current condition of its timing estimator. ■ **no-std-ext-time-src**: The AP using this profile has no standardized external time source. ■ **timestamp-offset-utc**: The AP has a timestamp offset based on UTC. ■ **reserved**: This setting is reserved for future use, and should not be used.
Time Error Value	The standard deviation of error in the time value estimate, in milliseconds. The default value is 0 milliseconds, and the supported range is 0-2,147,483,647 milliseconds.
P2P Device Management	Issue this command to advertise support for P2P device management. This setting is disabled by default.
P2P Cross Connect	Issue this command to advertise support for P2P Cross Connections. This setting is disabled by default.
Hotspot 2.0 Advertisement Protocol Type	The ANQP is used as the Hotspot 2.0 advertisement protocol types by default.
GAS comeback delay in milliseconds	At the end of the GAS comeback delay interval, the client may attempt to retrieve the query response using a Comeback Request Action frame. The supported range is 100-2000 milliseconds, and the default value is 500 milliseconds.
RADIUS Chargeable User Identity (RFC4372)	Include this parameter to enable the Chargeable-User-Identity RADIUS attribute defined by RFC 4372. Home networks can use this attribute to identify a user for the roaming transactions that take place outside of that home network.
RADIUS Location Data (RFC5580)	Include this parameter to enable the Location Data and Operator-Name RADIUS attributes defined by RFC 5580. The first ANQP Domain Name profile will be used as the Operator-Name value. Enabling this parameter allows the RADIUS server to use user location data.

Parameter	Description
Subscription Remediation Server URL	URL of the subscription remediation server.
Deauth Imminent Reason URL	URL that explains the reason for deauthentication of the station.
Re-auth delay	Time to wait for a deauthenticated device to re-authenticate to the AP. Time is measured in seconds.
Session Information URL	URL of the session information server.
802.11u QoS MAP DSCP Expectations	This field specifies QoS MAP DSCP Expectations. It supports 21 sets and is entered as <value>:<up> separated by ',' where value is 0 to 3F or FF and up is 0-7. For example, 35:02,16:06.
802.11u QoS MAP DSCP Ranges	This field specifies QoS MAP DSCP Ranges. It supports 8 sets and is entered as <low>:<high> separated by ',' where the value for low and high is 0 to 3F or FF.

The following CLI commands configure a hotspot 2.0 profile:

```
wlan hotspot h2-profile <profile-name>
   access-network-type emergency-services|personal-device|private|private-guest|public-
   chargeable|public-free|test|wildcard
   addtl-roam-cons-ois <addtl-roam-cons-ois>
   advertisement-profile <profile-name>
   advertisement-protocol anqp|eas|mih-cmd-event|mih-info|rsvd
   asra
   clone <profile-name>
   comeback-mode
   gas-comeback-delay
   grp-frame-block
   hessid <id>
   hotspot-enable
   internet
   no ..
   p2p-cross-connect
   p2p-dev-mgmt
   pame-bi
   query-response-length-limit <query-response-length-limit>
   radius_cui
   radius_loc_data
   roam-cons-len-1 0|3|5
   roam-cons-len-2 0|3|5
   roam-cons-len-3 0|3|5
   roam-cons-oi-1 <roam-cons-oi-1>
   roam-cons-oi-2 <roam-cons-oi-1>
   roam-cons-oi-3 <roam-cons-oi-1>
   time-advt-cap no-std-ext-timesrc|timestamp-offset-utc |reserved
   time-error <milliseconds>
   time-zone <time-zone>
   venue-group <venue-group>
   venue-type <venue-type>
```

Configuring Hotspot Advertisement Profiles

An advertisement profile defines a set of ANQP and H2QP profiles for the hotspot feature. Advertisement profiles can reference multiple instances of some ANQP and H2QP profile types, but only a single instance of other ANQP and H2QP profiles. The table below shows how the different ANQP and H2QP profile types can be associated to a single advertisement profile.

Table 210: *Hotspot Advertisement Profile Associations*

One Instance per Advertisement Profile	Multiple Instances per Advertisement Profile
■ ANQP IP address availability profile ■ H2QP WAN metrics profile ■ H2QP connection capability profile	■ ANQP venue name profile ■ ANQP network authentication profile ■ ANQP foaming consortium profile ■ ANQP NAI realm profile ■ ANQP 3GPP cellular network profile ■ H2QP operator friendly name profile ■ H2QP operating class indication profile ■ ANQP domain Name profile



For additional information on each of these profile types, see Hotspot 2.0 on page 1020

Configuring Advertisement Profile

The hotspot profile associates directly with the default advertisement profile. You can create separate ANQP and H2QP profiles under advertisement profile, and configure the required parameters for each profile. These advertisement profiles can then be associated to any of the hotspot profiles.

The following procedure describes how to configure an advertisement profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the **All Profiles** table, expand **Wireless LAN**.
3. Expand **Advertisement**. To create a new advertisement profile, click **+** in **Advertisement Profile: New Profile**.
4. Enter the **Profile Name** and **OSU SSID**.
5. Click **Submit**.
6. Select the created advertisement profile and configure the ANQP and H2QP profiles.

For details, refer to Hotspot 2.0.

Associating an Advertisement Profile to a Hotspot Profile

The following procedure describes how to associate an advertisement profile to a hotspot profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the **All profiles** table, expand **Wireless LAN.**
3. Expand **Hotspot 2.0**. Select a hotspot profile.
4. Select **Advertisement**.
5. Select an advertisement profile from the **Advertisement Profile** drop-down list.
6. Click **Submit**.
7. Click **Pending changes**.
8. In the **Pending changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure an advertisement profile:

```
wlan hotspot advertisement profile <profile-name>
   anqp-3gpp-nwk-profile <profile-name>
   anqp-domain-name-profile <profile-name>
   anqp-ip-addr-avail-profile <profile-name>
   anqp-nai-realm-profile <profile-name>
   anqp-nwk-auth-profile <profile-name>
   anqp-roam-cons-profile <profile-name>
   anqp-venue-name-profile <profile-name>
   clone <profile-name>
   h2qp-conn-cap-profile <profile-name>
   h2qp-op-cl-profile <profile-name>
   h2qp-operator-friendly-profile <profile-name>
   h2qp-wan-metrics-profile <profile-name>
   no ...
```

Associating Advertisement Profile to Hotspot 2.0 Profile

The settings in the ANQP and H2QP profiles referenced by the Advertisement profile will not be sent to clients until you associate the advertisement profile with an active hotspot 2.0 profile. By default, all hotspot 2.0 profiles reference the **default** advertisement profile.

The following procedure describes how to associate a different advertisement profile to a hotspot 2.0 profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration> System > Profiles**.
2. In the **All Profiles** table, expand **Wireless LAN**.
3. Expand **Hotspot 2.0**. The list of available hotspot 2.0 profiles appears in the **All Profiles** table.
4. Select an existing Hotspot 2.0 profile from the list of Hotspot 2.0 profiles.
5. Click **Advertisement** for the selected Hotspot 2.0 profile.
6. In the **All Profiles** table, click the **Advertisement Profile** drop-down list and select a different advertisement profile name.
7. Click **Submit**.

The following CLI commands associate a different advertisement profile to a hotspot 2.0 profile:

```
wlan hotspot hs2-profile <hotspot-profile-name>
   advertisement-profile <advertisement-profile-name>
```

Configuring ANQP Venue Name Profiles

Use this profile to define the venue group and venue type information which is sent in an ANQP information element in a GAS query response.

To send the values configured in this profile to clients, associate this profile with an advertisement profile, then associate the advertisement profile with a Hotspot 2.0 profile. For additional details, see Configuring Hotspot Advertisement Profiles on page 1026.

The following procedure describes how to configure an ANQP venue name profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the **All Profiles** table, expand **Wireless LAN**.
3. Expand **ANQP Venue Name**. The list of available ANQP venue name profiles appears in the **All Profiles** table.
4. Select an existing ANQP venue name profile from the list of profiles or create a new ANQP venue name profile by clicking **+**.
5. Configure the following parameters as desired, then click **Submit** to save your settings.

Table 211: *ANQP Venue Name Profile Parameters*

Parameter	Description
Venue Group	Specify one of the following venue groups to be advertised in the ANQP IEs from APs associated with this profile. The default setting is unspecified. ■ assembly ■ business ■ educational ■ factory-or-industrial ■ institutional ■ mercantile ■ outdoor ■ reserved ■ residential ■ storage ■ unspecified ■ Utility-Misc ■ Vehicular
Venue Language Code	An ISO 639 language code that identifies the language used in the Venue Name field.
Venue Name	Venue name to be advertised in the ANQP IEs from APs associated with this profile. If the venue name includes spaces, the name must be enclosed in quotation marks, e.g. “Midtown Shopping Center”.
Venue Type	Specify a venue type to be advertised in the IEs from APs associated with this hotspot profile. The complete list of supported venue types is described the table below
Venue Name in HEX	Network venue name in HEX.
Venue URL	The venue URL to access additional information about the venue.

The following CLI commands configure an ANQP venue name profile:

```
wlan hotspot anqp-venue-name-profile <profile-name>
   clone <profile-name>
   no ...
   venue-group outdoor|reserved|utility-
   misc|vehicular|assembly|business  educational|factory-or-
   industrial|institutional|mercantile|residential|   storage|unspecified
   venue-language <language>
   venue-name <venue-name>
   venue-type <venue-type>
```

Venue Types

The following list describes the different venue types that may be configured in a Hotspot 2.0 or ANQP Venue Name profile:

■ assembly-amphitheater ■ assembly-amusement-park ■ assembly-arena ■ assembly-bar	■ business-fire-station ■ business-police-station ■ business-post-office ■ business-professional-office ■ business-research-and-	■ mercantile-shopping-mall ■ outdoor-bus-stop ■ outdoor-city-park ■ outdoor-kiosk

■ assembly-coffee-shop ■ assembly-convention-center ■ assembly-emer-coord-center ■ assembly-library ■ assembly-museum ■ assembly-passenger-terminal ■ assembly-restaurant ■ assembly-stadium ■ assembly-theater ■ assembly-worship-place ■ assembly-zoo ■ business-attorney ■ business-bank ■ business-doctor	development ■ educational-primary-school ■ educational-secondary-school ■ educational-university ■ industrial-factory ■ institutional-alcohol-or-drug-rehab ■ institutional-group-home ■ institutional-hospital ■ institutional-prison ■ institutional-terminal-care ■ mercantile-automotive-service-station ■ mercantile-gas-station ■ mercantile-grocery ■ mercantile-retail	■ outdoor-muni-mesh-nwk ■ outdoor-rest-area ■ outdoor-traffic-control ■ residential-boarding-house ■ residential-dormitory ■ residential-hotel ■ residential-private-residence ■ unspecified ■ vehicular-airplane ■ vehicular-automobile ■ vehicular-bus ■ vehicular-ferry ■ vehicular-motor-bike ■ vehicular-ship ■ vehicular-train

Configuring ANQP Network Authentication Profiles

Use the ANQP Network Authentication profile to define the authentication type used by the Hotspot network.

To send the values configured in this profile to clients, associate this profile with an advertisement profile, then associate the advertisement profile with a Hotspot 2.0 profile. For additional details, see Configuring Hotspot Advertisement Profiles on page 1026.

The following procedure describes how to configure an ANQP network authentication profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the **All Profiles** table, expand **Wireless LAN**.
3. Expand **ANQP Network Authentication**. The list of available ANQP network authentication profiles appears in the **All Profiles** table.
4. Select an existing ANQP network authentication profile from the list of profiles or create a new profile by clicking **+**.
5. Configure the following parameters as desired, then click **Submit** to save your settings.

Table 212: *ANQP Network Authentication Profile Parameters*

Parameter	Description
Type of Network Authentication	Network Authentication Type being used by the hotspot network. ■ **acceptance**: Network requires the user to accept terms and conditions. This option requires you to specify a redirection URL string as an IP address, FQDN or URL. ■ **dns-redirection**: Additional information on the network is provided through DNS redirection. This option requires you to specify a redirection URL string as an IP address, FQDN or URL.

Parameter	Description
	■ **http-https-redirection**: Additional information on the network is provided through HTTP or HTTPS redirection. ■ **online-enroll**: Network supports online enrollment.
Network Authentication URL	URL, IP address, or FQDN used by the hotspot network for the **acceptance**, **dns-redirection**, or **online-enroll** network authentication types.

The following CLI commands configure an ANQP network authentication profile :

```
wlan hotspot anqp-nwk-auth-profile <profile-name>
   clone <profile-name>
   no ...
   nwk-auth-type acceptance|dns-redirection|http-https-redirection|online-enroll
   url <url>
```

Configuring ANQP Domain Name Profiles

This profile defines the hotspot operator domain name to be sent in an ANQP information element in a GAS query response.

To send the values configured in this profile to clients, associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For additional details, see Configuring Hotspot Advertisement Profiles on page 1026.

The following procedure describes how to onfigure an ANQP domain name profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the **All Profiles** table, expand **Wireless LAN**.
3. Expand **ANQP Domain Name.**
4. Select an existing profile from the list of profiles or create a new profile by clicking **+**.
5. In the **Domain Name** field, enter the domain name of the hotspot operator. Ensure that the alphanumeric text string is 255 characters or less.
6. Click **Save**.

The following CLI commands configure an ANQP domain name profile in config mode:

```
wlan hotspot anqp-domain-name-profile <profile-name>
   clone <profile-name>
   domain-name <domain-name>
   no ...
```

Configuring ANQP IP Address Availability Profiles

Use this profile to specify the types of IPv4 and IPv6 IP addresses available in the hotspot network. This information is sent in an ANQP information element in a GAS query response.

To send the values configured in this profile to clients, associate this profile with an advertisement profile, then associate the advertisement profile with a Hotspot 2.0 profile. For additional details, see Configuring Hotspot Advertisement Profiles on page 1026.

The following procedure describes how to configure an ANQP IP address availability profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the **All Profiles** table, expand **Wireless LAN**.

3. Expand **ANQP IP Address Availability**
4. Select an existing profile from the list of profiles or create a new profile by clicking **+**.
5. Configure the following parameters as desired, then click **Submit** to save your settings.

Table 213: *ANQP IP Address Availability Profile Parameters*

Parameter	Description
IPv4 Address Availability Type	Indicate the availability of an IPv4 network by clicking the **IPv4 Address Availability Type** drop-down list and selecting one of the following options: ■ availability-unknown: Network availability cannot be determined. ■ not-available: Network is not available. ■ port-restricted: Some ports are restricted (e.g., the network blocks port 110 to restrict POP mail). ■ port-restricted-double-nated: Some ports are restricted and multiple routers perform network address translation. ■ port-restricted-single-nated: Some ports are restricted and a single router performs network address translation. ■ private-double-nated: Network is a private network with multiple routers doing network address translation. ■ private-single-nated: Network is a private network a single router doing network address translation. ■ public: Network is a public network.
IPv6 Address Availability Type	Indicate the availability of an IPv6 network by clicking the **IPv6 Address Availability Type** drop-down list and selecting one of the following options: ■ available: An IPv6 network is available. ■ availability-unknown: Network availability cannot be determined. ■ not-available: Network is not available.

The following CLI commands configure ANQP IP address availability profile in config mode:

```
wlan hotspot anqp-ip-addr-avail-profile <profile-name>
   clone <profile-name>
   ipv4-addr-avail availability-unknown|not-available|port-restricted|port-restricted-
   double-nated|port-restricted-single-nated|private-double-nated|private-single-nated
   ipv6-addr-avail available|availability-unknown|not-available
   no ...
```

Configuring ANQP NAI Realm Profiles

An AP's NAI Realm profile identifies and describes a NAI realm accessible using the AP, and the method that this NAI realm uses for authentication. These settings configured in this profile determine the NAI realm elements that are included as part of a GAS Response frame.

To send the values configured in this profile to clients, associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For additional details, see Configuring Hotspot Advertisement Profiles on page 1026.

The following procedure describes how to configure an ANQP NAI Realm profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the **All Profiles** table, expand **Wireless LAN**.
3. Expand **ANQP NAI Realm**
4. Select an existing profile from the list of profiles or create a new profile by clicking **+** in the **+**.
5. Configure the following parameters as desired, then click **Submit** to save your settings.

Table 214: *ANQP NAI Realm Profile Parameters*

Parameter	Description
NAI Realm name	Name of the NAI realm. The realm name is often the domain name of the service provider.
NAI Realm Encoding	Issue this command if the NAI realm name is a UTF-8 formatted character string that is not formatted in accordance with IETF RFC 4282.
NAI Realm EAP Method 1/2/3/4	Select one of the options below to identify the EAP authentication method supported by the hotspot realm. ■ crypto-card: Crypto card authentication ■ eap-aka: EAP for UMTS Authentication and Key Agreement ■ eap-sim: EAP for GSM Subscriber Identity Modules ■ eap-tls: EAP-Transport Layer Security ■ eap-ttls: EAP-Tunneled Transport Layer Security ■ generic-token-card: EAP-GTC ■ identity: EAP Identity type ■ notification: The hotspot realm uses EAP Notification messages for authentication. ■ one-time-password: Authentication with a single-use password ■ peap: Protected Extensible Authentication Protocol ■ peap-mschapv2: Protected Extensible Authentication Protocol with Microsoft CHAP version 2
NAI Realm Authentication Param ID 1/2/3/4	Use the **NAI Realm Authentication Param ID** parameter to send the one of the following authentication methods for the primary NAI realm ID. ■ credential-type: The specified authentication ID uses credential authentication. ■ expanded-eap: The specified authentication ID uses the expanded EAP authentication method. ■ expanded-inner-eap: The specified authentication ID uses the expanded inner EAP authentication method. ■ inner-auth-eap: The specified authentication ID uses inner EAP authentication type. ■ non-eap-inner-auth: The specified authentication ID uses non-EAP inner authentication type.
NAI Realm Authentication Param Value 1/2/3/4	Use the **NAI Realm Authentication Param Value** parameter select an authentication value for the authentication method specified by the **NAI Realm Authentication Param ID** parameter. ■ cred-cert: Credential - Certificate ■ cred-hw-token: Credential - Hardware Token ■ cred-nfc: Credential - NFC ■ cred-none: Credential - None ■ cred-rsvd: Credential - Reserved ■ cred-sim: Credential - SIM ■ cred-soft-token: Credential - Soft Token ■ cred-user-pass: Credential - Username and password ■ cred-usim: Credential - USIM ■ cred-vendor-spec: Credential - Vendor-specific ■ eap-crypto-card: EAP Method - Crypto-card ■ eap-generic-token-card: EAP Method - Generic-Token-Card ■ eap-identity: EAP Method - Identity ■ eap-method-aka: EAP Method - AKA ■ eap-method-sim: EAP Method - SIM - GSM Subscriber Iden ■ eap-method-tls: EAP Method - TLS - Transport Layer Sec ■ eap-method-ttls: EAP Method - TTLS - Tunneled Transport Security ■ eap-notification: EAP Method - Notification

Parameter	Description
	■ eap-one-time-password: EAP Method - One-Time-Password ■ eap-peap: EAP Method - PEAP ■ eap-peap-mschapv2: EAP Method - PEAP MSCHAP V2 ■ non-eap-chap: Non-EAP Method - CHAP ■ non-eap-mschap: Non-EAP Method - MSCHAP ■ non-eap-mschapv2: Non-EAP Method - MSCHAPv2 ■ non-eap-pap: Non-EAP Method - PAP ■ non-eap-rsvd: Non-EAP Method - Reserved for future use ■ reserved: Reserved for Future use
NAI Home Realm	Mark the realm in this profile as the NAI Home Realm.

The following CLI commands configure an ANQP NAI realm profile:

```
wlan hotspot anqp-nai-realm-profile <profile-name>
   clone <profile-name>
   nai-home-realm
   nai-realm-auth-id-1|nai-realm-auth-id-2 {credential-type|expanded-eap|expanded-inner-
   eap|inner-auth-eap|non-eap-inner-auth|tunneled-eap-credential-type}
   nai-realm-auth-value-1|nai-realm-auth-value-2 {cred-cert|cred-hw-token|cred-nfc|cred-
   none|cred-rsvd|cred-sim|cred-soft-token|cred-user-pass|cred-usim|cred-vendor-spec|eap-
   crypto-card|eap-generic-token-card|eap-identity|eap-method-aka|eap-method-sim|eap-
   method-tls|eap-method-ttls|eap-notification|eap-one-time-password|eap-peap|eap-peap-
   mschapv2|non-eap-chap|non-eap-mschap|non-eap-mschapv2|non-eap-pap|non-eap-
   rsvd|reserved}
   nai-realm-eap-method crypto-card|eap-aka|eap-sim|eap-tls|eap-ttls|generic-token-
     card|identity|notification|one-time-password|peap|peap-mschapv2
   nai-realm-encoding
   nai-realm-name <nai-realm-name>
   no ...
```

Configuring ANQP Roaming Consortium Profiles

OIDs are assigned to service providers when they register with the IEEE registration authority. You can specify the OI for the hotspot's service provider in the ANQP Roaming Consortium profile using the ANQP Roaming Consortium Profile. The Hotspot 2.0 profile also allows you to define and send up to three additional roaming consortium OIs for the service provider's top three roaming partners.

To send the values configured in this profile to clients, associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For additional details, see Configuring Hotspot Advertisement Profiles on page 1026.



Ensure that the Hotspot 2.0 profile Additional Roaming OI Consortium number is re-visited each time you add or remove one of those profile.

The following procedure describes how to configure an ANQP roaming consortium profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the **All Profiles** table, expand **Wireless LAN**.
3. Expand **ANQP Roaming Consortium.**
4. Select an existing profile from the list of profiles or create a new profile by clicking **+**.
5. Configure the following parameters as desired, then click **Submit** to save your settings.

Table 215: *ANQP Roaming Consortium Profile Parameters*

Parameter	Description
Roaming consortium OI Len	Length of the OI. The value of the **Roaming consortium OI Len** parameter must equal upon the number of octets of the **Roaming Consortium OI** field. ■ **0**: 0 Octets in the OI (Null) ■ **3**: OI length is 24-bit (3 Octets) ■ **5**: OI length is 36-bit (5 Octets)
Roaming Consortium OI	Send the specified roaming consortium OI in a GAS query response. The OI must be a hexadecimal number 3-5 octets in length.

The following CLI commands configure an ANQP roaming consortium profile from the CLI:

```
wlan hotspot anqp-roam-cons-profile <profile-name>
   clone <profile-name>
   no ...
   roam-cons-oi <roam-cons-oi>
   roam-cons-oi-len <roam-cons-oi-len>
```

Configuring ANQP 3GPP Cellular Network Profiles

Use this profile to define priority information for a 3GPP Cellular Network used by hotspots that have roaming relationships with cellular operators.

To send the values configured in this profile to clients, associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For additional details, see Configuring Hotspot Advertisement Profiles on page 1026.

The following procedure describes how to configure an ANQP 3GPP cellular network profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the **All Profiles** table, expand **Wireless LAN**.
3. Expand **ANQP 3GPP Cellular Network.**
4. Select an existing profile from the list of profiles or create a new profile by clicking **+**.
5. Enter a **Profile Name** and configure the following parameters as desired, then click **Submit** to save your settings.

Table 216: *ANQP 3GPP Cellular Network Profile Parameters*

Parameter	Description
3GPP PLMN1	The PLMN value of the highest-priority network. The PLMN is comprised of a 12-bit Mobile Country Code and the 12-bit Mobile Network Code.
3GPP PLMN2	The PLMN value of the second-highest priority network. The PLMN is comprised of a 12-bit Mobile Country Code and the 12-bit Mobile Network Code.
3GPP PLMN3	The PLMN value of the third-highest priority network. The PLMN is comprised of a 12-bit Mobile Country Code and the 12-bit Mobile Network Code.
3GPP PLMN4	The PLMN value of the fourth-highest priority network. The PLMN is comprised of a 12-bit Mobile Country Code and the 12-bit Mobile Network Code.

Parameter	Description
3GPP PLMN5	The PLMN value of the fifth-highest priority network. The PLMN is comprised of a 12-bit Mobile Country Code and the 12-bit Mobile Network Code.
3GPP PLMN6	The PLMN value of the sixth-highest priority network. The PLMN is comprised of a 12-bit Mobile Country Code and the 12-bit Mobile Network Code.

The following CLI commands configure an ANQP 3GPP network profile in config mode:

```
wlan hotspot anqp-3gpp-nwk-profile <profile-name>
   3gpp_plmn1    <3GPP PLMN1 data>
   3gpp_plmn2    <3GPP PLMN2 data>
   3gpp_plmn3    <3GPP PLMN3 data>
   3gpp_plmn4    <3GPP PLMN4 data>
   3gpp_plmn5    <3GPP PLMN5 data>
   3gpp_plmn6    <3GPP PLMN6 data>
   clone <profile-name>
   enable
   no ...
```

Configuring H2QP Connection Capability Profiles

Use this profile to specify hotspot protocol and port capabilities. This information is sent in a ANQP information element in a GAS query response.

To send the values configured in this profile to clients, associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For additional details, see Configuring Hotspot Advertisement Profiles on page 1026.

The following procedure describes how to configure a H2QP connection capability profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the All profiles table, expand **Wireless LAN**.
3. Expand **H2QP Connection Capability.**
4. Select an existing profile from the list of profiles or create a new profile by clicking **+**.
5. Configure the following parameters as desired, then click **Submit** to save your settings.

Table 217: *ANQP Connection Capability Profile Parameters*

Parameter	Description
H2QP Connection Capability ICMP Port	Select this option to enable the ICMP port. (port 0)
H2QP Connection Capability FTP port (TCP Protocol)	Select this option to enable the FTP port. (port 20)
H2QP Connection Capability SSH port (TCP Protocol)	Select this option to enable the SSH port. (port 22)
H2QP Connection Capability HTTP port (TCP Protocol)	Select this option to enable the HTTP port. (port 80)
H2QP Connection Capability TLS VPN port (TCP Protocol)	H2QP Connection Capability TLS VPN port (TCP Protocol).

Parameter	Description
H2QP Connection Capability PPTP VPN port (TCP Protocol)	Select this option to enable the PPTP port used by IPsec VPNs (port 1723).
H2QP Connection Capability VOIP port (TCP Protocol)	Select this option to enable the TCP VoIP port (port 5060).
H2QP Connection Capability VOIP port (UDP Protocol)	Select this option to enable the UDP VoIP port. (port 5060)
H2QP Connection Capability IKEv2 port for IPSec VPN	Select this option to enable the IPsec VPN port. (ports 500, 4500 and 0)
H2QP Connection Capability May be used by IKEv2 port for IPSec VPN	Select this option to enable the IKEv2 port 4500.
H2QP Connection Capability ESP port (Used by IPSec VPN)	Include this parameter to enable the Encapsulating Security Payload port used by IPsec VPNs. (port 0)

The following CLI commands configure a H2QP connection capability profile:

```
wlan hotspot h2qp-conn-capability-profile <profile>
   clone <profile-name>
   esp
   icmp
   no ...
   tcp-ftp
   tcp-http
   tcp-pptp-vpn
   tcp-ssh
   tcp-tls-vpn
   tcp-voip
   udp-ike2-4500
   udp-ike2-500
   udp-ipsec-vpn
   udp-voip
```

Configuring H2QP Operator Friendly Name Profiles

This profile defines an operator-friendly name sent by devices using this profile.

To send the values configured in this profile to clients, associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For additional details, see Configuring Hotspot Advertisement Profiles on page 1026.

The following procedure describes how to configure a H2QP operating class profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the **All Profiles** table, expand **Wireless LAN**.
3. Expand **H2QP Operator Friendly Name.**
4. Select an existing profile from the list of profiles or create a new profile by clicking **+**.
5. Configure the following parameters as desired, then click **Submit** to save your settings.

Table 218: *H2QP Operator Friendly Name Profile Parameters*

Parameter	Description
Operator Friendly Name Language Code	An ISO 639 language code that identifies the language used in the **Operator Friendly Name** field
Operator Friendly Name	An operator-friendly name sent by devices using this profile. The name can be up to 64 alphanumeric characters, and can include special characters and spaces. If the name includes quotation marks (“), include a backslash character (\) before each quotation mark. (e.g. \"example\")
Operator Friendly Name in HEX (no separators)	Operator friendly name in HEX.

The following CLI commands configure a H2QP operator friendly name profile:

```
wlan hotspot h2qp-operator-friendly-name-profile <profile>
   clone <profile-name>
   no ...
   op-fr-name <op-fr-name>
   op-lang-code <op-lang-code>
```

Configuring H2QP Operating Class Indication Profiles

The values configured in this H2QP Operating Class Indication profile list the channels on which the hotspot is capable of operating. It may be useful where, for instance, a mobile device discovers a hotspot in the 2.4 GHz band but finds it is dual-band and prefers the 5 GHz band.

To send the values configured in this profile to clients, associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For additional details, see Configuring Hotspot Advertisement Profiles on page 1026.

The following procedure describes how to configure a H2QP operating class indication profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the **All Profiles** table, expand **Wireless LAN**.
3. Expand **H2QP Operating Class Indication.**
4. Select an existing profile from the list of profiles or create a new profile by clicking **+**.
5. In the **H2QP Operating Class** field, enter a valid operating class value. (For a definition of these global operating classes refer to Table E-4 of IEEE Std 802.11-2012, Annex E.)
6. Click **Submit**.

The following CLI commands configure a H2QP operating class profile:

```
wlan hotspot h2qp-op-cl-profile <profile>
   clone <profile-name>
   op-cl <1-255>
```

Configuring H2QP WAN Metrics Profiles

Use this profile to specify the WAN status and link metrics for your hotspot.

To send the values configured in this profile to clients, associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For additional details, see Configuring Hotspot Advertisement Profiles on page 1026.

The following procedure describes how to configure an ANQP venue name profile:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In the **All Profiles** table, expand **Wireless LAN**.
3. Expand **H2QP WAN Metrics.**
4. Select an existing profile from the list of profiles or create a new profile by clicking **+** in the **New Profile** table.
5. Configure the following parameters as desired, then click **Submit** to save your settings.

Table 219: *H2QP WAN Metrics Profile Parameters*

Parameter	Description
H2QP WAN metrics link status	Define the status of the WAN Link by clicking the **H2QP WAN metrics link status** drop-down list, and selecting one of the following values. The default link status is **reserved**, which indicates that the link status is unknown or unspecified ■ **link down:** WAN link is down. ■ **link test**: WAN link is currently in a test state. ■ **link up**: WAN link is up. ■ **reserved**: This parameter is reserved by the Hotspot 2.0 specification, and cannot be configured. Default: reserved
H2QP WAN metrics symmetric WAN link	Select this check box to indicate that the WAN Link has same speed in both the uplink and downlink directions.
H2QP WAN metrics link at capacity	Select this check box to indicate that the WAN Link has reached its maximum capacity. If this parameter is enabled, no additional mobile devices will be permitted to associate to the hotspot AP.
WAN Metrics uplink speed	This parameter defines the current WAN uplink speed in Kbps. If no value is set, this parameter will show a default value of 0 to indicate that the uplink speed is unknown or unspecified. Range: 0 - 2147483647, Default: 0
WAN Metrics downlink speed	This parameter defines the current WAN backhaul downlink speed in Kbps. If no value is set, this parameter will show a default value of 0 to indicate that the downlink speed is unknown or unspecified. Range: 0 - 2147483647, Default: 0
WAN Metrics uplink load	This parameter defines the percentage of the WAN uplink that is currently utilized. If no value is set, this parameter will show a default value of 0 to indicate that the downlink speed is unknown or unspecified. Range: 0-100; Default: 0
WAN Metrics downlink load	This parameter defines the percentage of the WAN downlink that is currently in use. If no value is set, this parameter will show a default value of 0 to indicate that the downlink speed is unknown or unspecified. Range: 0-100; Default: 0
WAN Metrics load measurement duration	Duration over which the downlink load is measured, in tenths of a second. Range: 0-65535; Default: 0

The following CLI commands configure a H2QP WAN metrics profile:

```
wlan hotspot h2qp-wan-metrics-profile <profile-name>
   at-capacity
   clone <profile-name>
   downlink-load
   downlink-speed
   load-dur
   no ...
```

```
symm-link
uplink-load
uplink-speed
wan-metrics-link-status link_down|link_test|link_up|reserved
```

The SDN Controller provides an improved networking infrastructure through the following enhancements:

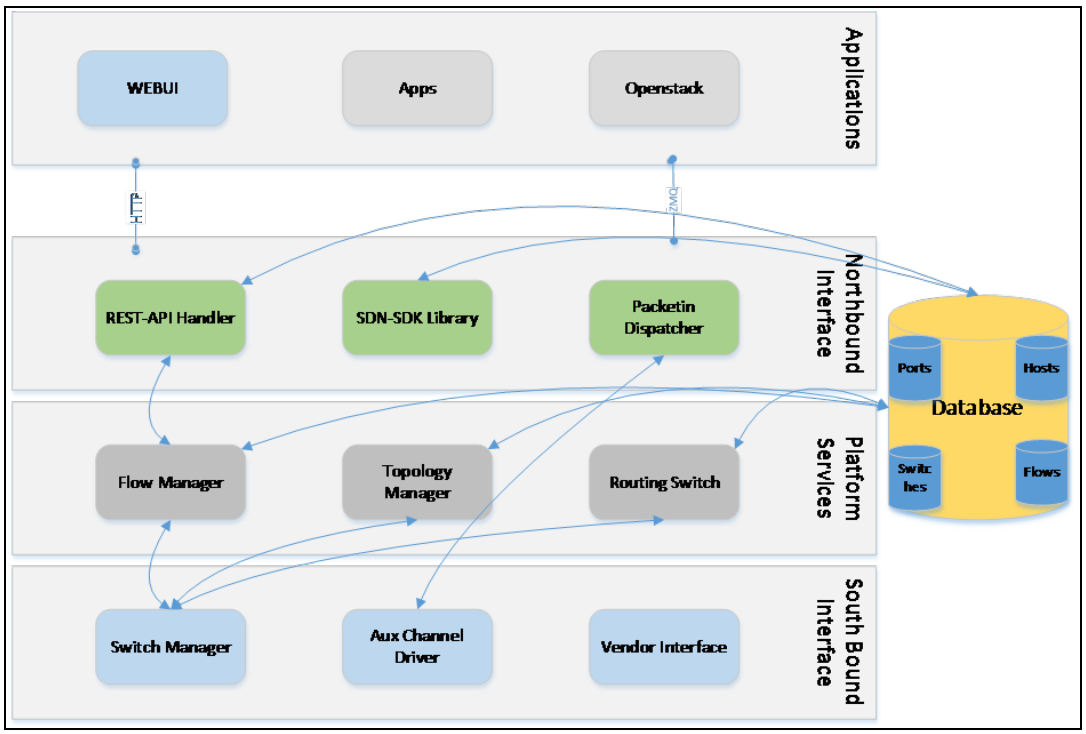
- Separation of control-plane and data-plane functions
- Centralized manageability
- Dynamic programmability of network devices

Traditional networks can experience high latency and inflexibility, as the number of new applications and features continues to grow. All control and forwarding functions take place on the same device, and features can only be provisioned statically through manual intervention. The SDN Controller provides a more efficient and simple way to build, deliver, and manage features throughout the network.

This section describes the following modules that form the SDN Controller:

- Southbound Interface on page 1041
- SDN Controller Configuration on Mobility Master on page 1042
- SDN Platform Services on page 1042
- Northbound API on page 1051

Figure 122 *SDN Controller Architecture*



Southbound Interface

The Southbound Interface is a collection of drivers that handles communication to all data-plane elements in the network.

OpenFlow Driver

The OpenFlow Driver supports the dynamic manipulation of network devices and the separation of fast-packet forwarding (data-plane) from high-level routing (control-plane). Data-plane functions reside on the data-plane element (switch), while control-plane functions have migrated to a separate controller. The OpenFlow switch and controller communicate through the OpenFlow protocol to provide functions such as host discovery and packet handling. OpenFlow allows users to run and manage multiple instances of the control-plane and data-plane from a centralized location.



The SDN Controller supports OpenFlow versions 1.0 and 1.3.

Auxiliary Channel Driver

The Auxiliary Channel Driver carries all non-control data from the data-plane elements to the Mobility Master through UDP-based auxiliary channel connections. The auxiliary channel reduces bandwidth consumption and latency on the main channel, which must be used for critical functions such as flow programming or network state changes. All received data is sent to the subscribed Northbound APIs, which process and share the information with northbound applications.

The following items must be in place before an auxiliary channel connection can be used on the SDN Controller:

- Global OpenFlow configuration must be enabled on Mobility Master.
- The main channel must be UP.
- The listening port for the auxiliary connection must be configured (default connection is 6633).
- The source IP of the incoming UDP packets from the switch must be the same source IP used by the main channel connection.

For more information on configuring auxiliary channel ports, see Enabling SDN Controller on Mobility Master on page 1064

SDN Controller Configuration on Mobility Master

OpenFlow is an open communications interface between control plane and the forwarding layers of a network. OpenFlow allows dynamic manipulation of the forwarding plane for switches and routers. SDN architecture uses OpenFlow to enable software programs to manipulate the flow of packets in the network and to manage traffic based on the application's requirement.

OpenFlow Protocol v1.3 is used to achieve SDN with ArubaOS 8.0.0.0. An SDN controller runs on Mobility Master while an OpenFlow agent runs on the managed devices. For more information on OpenFlow agent, see OpenFlow Agent on page 1064. Mobility Master and the managed devices communicate over OpenFlow channels. The applications running in Mobility Master get all mDNS or SSDP packets seen by control plane on the managed devices. All outgoing mDNS or SSDP packets are originated by the application on Mobility Master.

For more information on SDN Controller configuration on Mobility Master, see Enabling SDN Controller on Mobility Master on page 1064

SDN Platform Services

SDN Platform Services gather and build the information required for core controller functions, including the following:

- Discovery of OpenFlow-capable devices and ports
- Discovery of all hosts and clients
- Discovery of the network topology
- Basic switching and routing
- Flow and policy programming
- Functionality to provide network paths between hosts
- Packet transmission
- Asynchronous event or state updates to Northbound applications

SDN Controller functions are achieved through the following services:

- Switch Discovery
- Topology discovery
- Host discovery
- Flow Management
- Packet Handling
- IPv6
- High Availability



The SDN Controller supports IPv6 flows and hosts. IPv6 host addresses can be learned in addition to IPv4 addresses.

Switch Discovery

Data-plane elements connect to the SDN Controller using the TCP. Each TCP connection is terminated by a switch manager and initiates OpenFlow messages between the data-plane element and the SDN Controller. When a new switch is discovered, the OpenFlow driver sends a message to the topology manager, which creates a new switch entry in the Switch Database. Applications can subscribe to this message type to receive a notification each time a switch is discovered.

Topology Discovery

The network topology displays the arrangement of switches and inter-switch links within the network. The complete topology view allows applications to:

- Make forwarding decisions
- Establish the shortest paths for flows
- Find alternate paths when links are congested or down
- Send low-latency flows on fast links and low-priority flows on slow links

The SDN Controller uses Layer-3 LLDP to discover links between switches. The LLDP frame is encapsulated in an IP packet with a unique source and destination IP address, creating a clear separation between standard LLDP packets and LLDP packets generated by the SDN Controller.

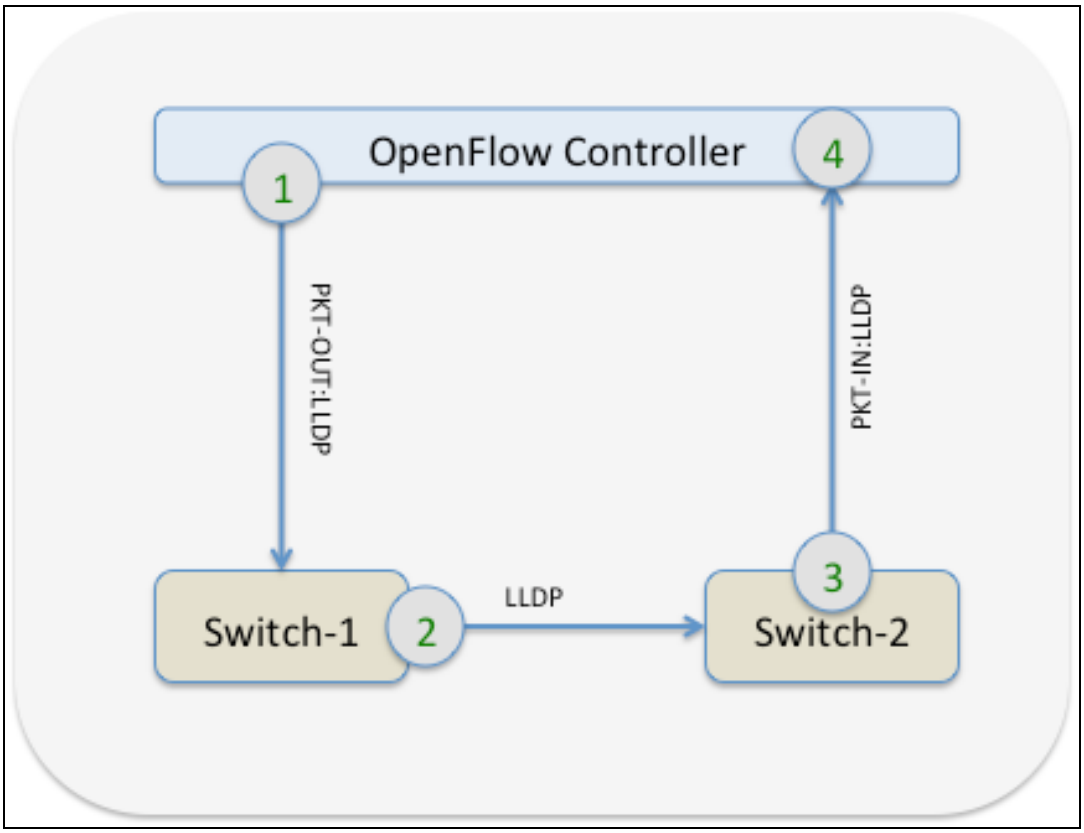
The following procedure generates the network topology:

1. For every link with an UP status, the controller sends an OpenFlow packet-out message with an LLDP frame to the source switch (switch-1).
2. Switch-1 sends the LLDP frame to switch-2 through the specified link.
3. Switch-2 sends the LLDP frame back to the SDN controller through a packet-in message.
4. Upon receiving the packet-in message, the controller creates a link between the port number (port) and datapath ID combination on switch-1 and port and datapath ID on switch-2.



This process only creates a link in one direction. Steps 1-4 must be repeated to create a link from the opposite direction

Figure 123 *Topology Discovery*



Host Discovery

The network topology can only be completed when all hosts are discovered by the controller. Hosts are defined by the port and datapath ID of the data-plane element to which they are connected. data-plane elements can run under hybrid mode (passive mode) or true OpenFlow mode (active mode).

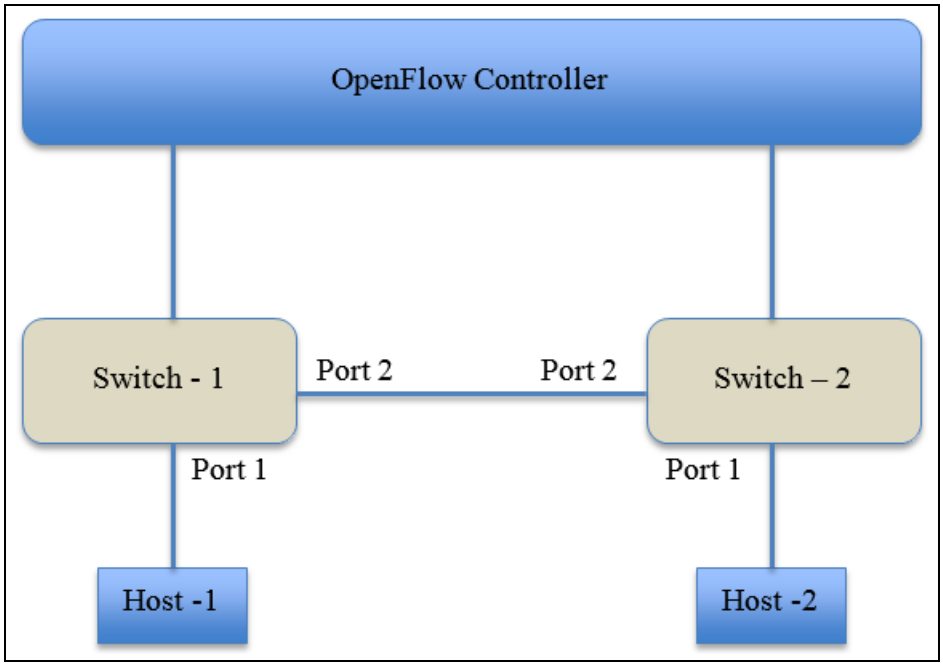
When a host connects to a data-plane element, the host generates ARP packets that contain important mapping and identification information. When the data-plane element registers to the controller, the data-plane element mirrors these ARP packets from the host to the controller through a packet-in message, in which the controller learns the IP-MAC binding, attachment point (port and datapath ID), and classification of the host.

Hosts can be classified as wireless or non-wireless, depending on the point of attachment. For example, if the port on the data-plane element is wireless, the host is marked as wireless.

Hosts can only be discovered by packets received through leaf links. Packets that are received through inter-switch links do not trigger host discovery on the controller.

Hosts can only be classified as wireless if all wireless tunnels contain a *bss* keyword.

Figure 124 *Host Discovery*



For example, in Figure 124, the ARP packet generated by Host-1 triggers a packet-in message from Switch-1 to the controller. Switch-2 also generates a packet-in message when the ARP packet is forwarded from Port 2 of Switch-1. However, the controller ignores the message from Switch-2 since the ARP packet is received through an inter-switch link (see Topology Discovery for more details on inter-switch links). The controller learns that Host-1 is connected to Port-1 of datapath ID Switch-1.

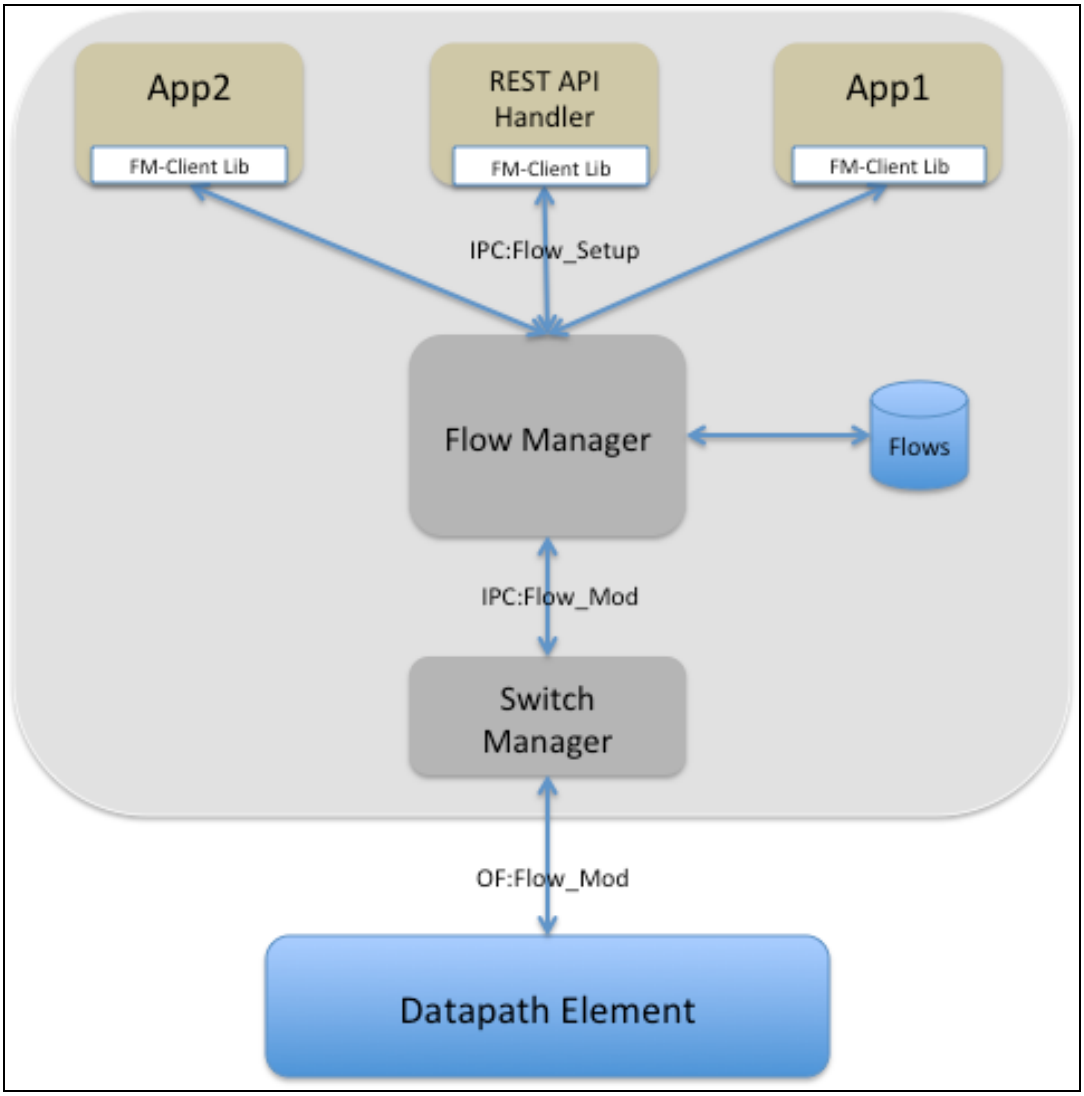
IPv6 hosts are also discovered similarly using ICMP Version 6 neighbor discovery and solicitation packets.

Host information is maintained in the Host Database, which is available to all SDN Controller applications and certain Northbound APIs. Host entries are aged out and deleted from the Host Database if no ARP packets for that host are received within the specified timeout period (default of 300 seconds).

Flow Management

The flow manager programs and maintains flows that are pushed by applications on the controller. APIs from the SDN Software Development Kit library interact with the flow manager to set up, update, or delete flows between applications and data-plane elements. A collection of flows that achieves a specific end-to-end policy or traffic forwarding process is referred to as a flow-group.

Figure 125 *Flow Push*



The flow manager includes the following functions:

Flow Installation

Flow installation is initiated when an application sends a flow setup message to the flow manager. The setup message contains the flow definition and allows the flow manager to check for any conflicting entries in the Flow Database. The flow is assigned a unique flow ID and flow-group ID by the application. The flow manager completes the installation process by sending a flow modification message to the switch manager that handles the target data-plane element.

Flow Match and Actions

The SDN Controller presents a uniform interface for flow installation, regardless of the OpenFlow version. Fields from incoming packets are matched against flow entries (match fields) to perform a specific set of actions. Refer to Table 220 to view the complete list of supported match fields and Table 221 to view the complete lists of supported actions.

Table 220: *Flow Match Fields*

Match Field	Description	Mand-atory	Mask-able	Type	Example
switch	DIPID of the switch.	Yes	No	String	00:00:00:1a:1e:00:3b:40
priority	Priority of the flow.	Yes	No	Integer	32768
idle-timeout	Idle-timeout, after which a flow is deleted based on the time period since activity was last detected.	No	No	Integer	30 seconds
hard-timeout	Hard-timeout, after which a flow is deleted based on the time period since the flow was created.	No	No	Integer	30 seconds
ingress-port	Ingress port of the packet.	No	No	Integer	1
src-mac	Source MAC address on the ether header.	No	No	String	00:1a:1e:00:3b:40
dst-mac	Destination MAC address on the ether header.	No	No	String	00:1a:1e:00:3b:40
ether-type	Type of ether header.	No	No	Integer	2048 (IPv4)
vlan	VLAN ID.	No	No	Integer	20
vlan-priority	VLAN priority.	No	No	Integer	6
src-ip	Source IP address on the IP header.	No	Yes	String	10.10.10.10
src-ip-mask	Number of bits to mask from the LSB.	No	—	Integer	8 (/24 mask)
dst-ip	Destination IP address on the IP header.	No	Yes	String	10.10.10.11
dst-ip-mask	Number of bits to mask from the LSB.	No	—	Integer	24 (/8 mask)
src-ipv6	Source IPv6 address on the IP header.	No	Yes	String	fe80::1a:1e0f:ff00:3b41/64
src-ipv6-mask	Number of bits to mask from the LSB.	No	—	Integer	64
dst-ipv6	Destination IPv6 address on the IP header.	No	Yes	String	fe80::1a:1e0f:ff00:3b41/64
dst-ipv6-mask	Number of bits to mask from the LSB.	No	—	Integer	32
icmpv6-type	ICMP version 6 type.	No	No	Integer	135 (neighbor

Match Field	Description	Mand-atory	Mask-able	Type	Example
					solicitation)
icmpv6-code	ICMPv6 code.	No	No	Integer	0
ip-tos	ToS bits on the IP header.	No	No	Integer	34
protocol	Protocol on the IP header.	No	No	Integer	6 (TCP)
src-port	Source port on the IP header.	No	No	Integer	5353 (MDNS)
dst-port	Destination port on the IP header.	No	No	Integer	5353 (MDNS)
app-name	Name of the application installing the flow.	No	No	String	airgroup
src-port-start	Start of the source port range.	No	No	Integer	5000
src-port-end	End of the source port range.	No	No	Integer	5010
dst-port-start	Start of the destination port range.	No	No	Integer	6000
dst-port-end	End of the destination port range.	No	No	Integer	6010

Table 221: *Flow Actions*

Action Field	Description	Example
output	Port on which the packet is sent out.	Controller Flood All Normal
set-vlan-id	Configures the VLAN ID on the VLAN header.	20
set-vlan-priority	Configures the VLAN priority on the VLAN header.	7
set-src-mac	Configures the source MAC address on the ether header.	00:1a:1e:00:3b:40
set-dst-mac	Configures the destination MAC address on the ether header.	00:1a:1e:00:3b:40
set-src-ip	Configures the source IP address on the IP header.	20.20.20.20
set-dst-ip	Configures the destination IP address on the IP header.	20.20.20.20
set-tos-bits	Configures ToS bits on the IP header.	42
set-flag	Updates session flags in the datapath to further process traffic. This action is only available on Aruba OpenFlow switches.	VH

Action Field	Description	Example
write-flag	Overwrites datapath session flags. This action is only available on Aruba OpenFlow switches.	VH
aruba-output	Sets a maximum packet number for the specified flow output. Only the specified number of packets is mirrored to the output for the flow. This action is only available on Aruba OpenFlow switches.	controller:10 (indicates that only the first 10 packets are forwarded to the controller)
set-appid	Sets an application ID on a datapath session after DPI has been performed. This action is only available on Aruba OpenFlow switches.	Netflix

Flow Update

Applications can use this function to update the action list associated with an existing flow. The application must obtain the new action list and the ID of the existing flow or flow-group to create a flow update request. After the request is accepted, the flow manager locates and updates the existing flow in the Flow Database.

Flow Deletion

The SDN Controller supports the following methods to delete a flow:

- Applications can explicitly call for a flow delete to remove a flow from the switch.
- Flows can be removed asynchronously through an idle-timeout or hard-timeout. Idle-timeout specifies the time period since activity was last detected for a flow. Hard-timeout specifies the time period since the flow was created.

After a flow is deleted, the controller removes all references to the flow and notifies northbound applications about the deletion.



Individual flows within a flow-group cannot be deleted; the entire flow-group must be deleted.

Flow Statistics

The flow manager sends statistics request messages to the data-plane elements to update statistics in the Flow Database every 30 seconds. Separate request messages are sent for each flow bucket, based on the corresponding cookie ID and cookie mask value. This bucket-based statistics collection improves the overall performance of the flow manager and data-plane elements since flow statistics can be processed through multiple (smaller) requests.

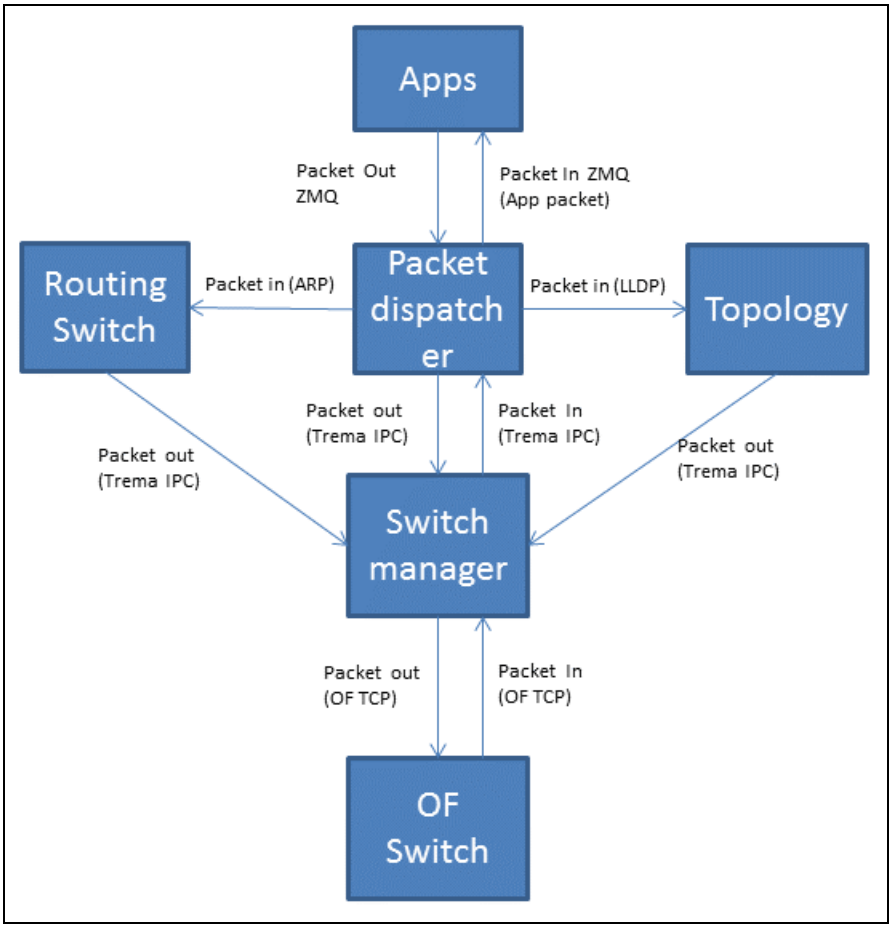
Databases

A copy of every network flow is maintained and readily accessible in the Flow Database. The storage of flows in a database prevents data or state loss during process crashes and provides information to read-only applications without requiring communication with any infrastructure processes.

Packet Handling

One of the major functions of the SDN Controller is to facilitate packet processing throughout the network. Packets that are sent from switches are called packet-in. Packets that are sent to switches are called packet-out. The following figure displays the packet-in and packet-out flow within the network.

Figure 126 *Packet Handling*



When a northbound application, routing switch, or topology manager sends a packet-out message, the packet is sent to the packetin-dispatcher, which handles all packet-related functionality, such as the port number and datapath ID to which the packet must be sent. The packet is then sent to the switch manager, which processes and sends the packet to the respective data-plane element (based on the designated port and datapath ID), where additional forwarding processes can be handled.

When a switch sends a packet-in message, the packet is sent to the packetin-dispatcher through the switch manager. The packetin-dispatcher classifies and delivers the packet to the routing switch, topology manager, or Northbound APIs.

OpenFlow Version Support

The SDN Controller supports OpenFlow versions 1.0 and 1.3. The controller dynamically negotiates with the DPEs to select the highest common version.

IPv6 Support

The SDN Controller supports IPv6 flows and hosts. Flow match conditions can include the following:

- IPv6 as an ether type value
- IPv6 address as an src-ip or dst-ip value
- ICMPv6 type and code

The controller can learn IPv6 addresses in addition to IPv4 addresses.

High Availability

With the centralization of features and management functions on the SDN Controller, a single instance of the controller creates a single point of failure. The SDN Controller supports the VRRP to reduce downtime and client traffic disruptions during network upgrades or unexpected failures. See Increasing Network Uptime With Redundancy Services for more details on high availability and VRRP.

VRRP

VRRP provides a redundancy solution, in which two or more systems, such as a primary and backup controller, share a virtual IP address. When the primary SDN Controller fails, the agents (DPEs) that are connected to the primary controller time out and reconnect to the backup controller associated with the same virtual IP.

Under the SDN Controller, the primary and backup controllers do not undergo a state sync. The following must be rebuilt on the backup controller:

- The backup controller must learn and compute the network topology.
- Previous hosts may not be available immediately on the system. The Host Database must be rebuilt.
- Upon switch discovery, applications must re-push all flows to the switches.

VRRP failovers are transparent for northbound applications that use REST APIs, as REST APIs continue to provide services despite any system failures. However, applications must re-subscribe to the ZeroMQ APIs.

Northbound API

The Northbound API makes the information built from the SDN Controller available for applications. The Northbound API consists of the following API types:

- **Synchronous**: Synchronous APIs are initiated by the client, and the information is presented through the server in response to the API. Synchronous APIs can be classified into two categories:
  - **Get and Fetch**: These APIs obtain information about the network without affecting the state of the network.
  - **Push, Post, and Modify**: These APIs modify the state of the network.
- **Asynchronous**: Asynchronous APIs notify northbound applications about changes in the network through server-to-client communication.

Synchronous APIs

Synchronous APIs are implemented through the Representational State Transfer API using standard GET HTTP, POST, and DELETE methods. Representational State Transfer provides a uniform interface between clients and servers, while allowing them to exist independently without any state transfers. A new request must be made by the client through a fetch mechanism or a single API each time the information is desired. Refer to the sections below to view the Representational State Transfer APIs that are available on the SDN Controller.

The SDN Controller provides multi-version support and backwards compatibility for Representational State Transfer APIs.

Switch API

The Switch API returns information about switches in the network, using the GET HTTP method.

Table 222: *Switch API Query Parameters*

Query Parameter	Definition
switch	Lists all switches in the network.
switch?dpid	Lists all switches associated with a specific datapath ID.

The output for this message type displays the following information:

Table 223: *Switch API Output Parameters*

Output Parameter	Definition
add-time	Time that the switch is added to the network.
auxiliary-id	ID number identifying the auxiliary channel.
auxiliary-status	Indicates if the auxiliary channel is **Up** or **Down**.
auxiliary-update-time	Time that the auxiliary channel is updated.
capabilities	Switch capabilities, such as flow statistics or port statistics.
description	Information about the switch, including the model and version of the switch.
disconnect-time	Time that the switch disconnects from the network.
ip	IP address of the switch.
port	Remote TCP port of the switch.
ports	Displays the list of ports on the switch.
name	Name of the port.
port-mac	MAC address of the port.
port-no	Port number.
rx-packets	Total number of packets received on the port.
status	Status of the port.
tx-packets	Total number of 802.11 packets transmitted by the port.
reconnect-time	Time that the switch reconnects to the network.
secure-connection	Indicates if a secure connection is **Enabled** or **Disabled** on the switch.
status	Indicates if the switch is **Up** or **Down**.
switch	MAC address of the switch.
version	OpenFlow version of the switch: ▪ v1.0 ▪ v1.3

The Switch API displays output similar to the example below:

```
#curl --insecure -b aruba-cookie https://10.4.251.105:4343/sdn/v1/switch
[
{
"add-time": "Mon Jun 29 07:54:10 2015\n",
"auxiliary-id": 1,
         "auxiliary-status": "Up",
"auxiliary-update-time": "Wed May 18 02:54:49 2016\n",
"capabilities": [
"Flow statistics",
"Table statistics",
"Port statistics",
"Queue statistics"
],
"description": "Aruba Networks, Inc. Aruba7210 VERSION 6.4 None None",
"disconnect-time": "Mon Jun 29 09:48:21 2015\n",
"ip": "10.4.251.79",
"port": 51898,
"ports": [
{
"name": "GE0/0/2",
"port-mac": "00:1a:1e:00:3b:43",
"port-no": 1,
"rx-packets": 0,
"status": 0,
"tx-packets": 0
},
{
"name": "GE0/0/3",
"port-mac": "00:1a:1e:00:3b:44",
"port-no": 2,
"rx-packets": 0,
"status": 0,
"tx-packets": 0
}
],
"reconnect-time": "Mon Jun 29 09:48:24 2015\n",
"secure-connection": "Disabled",
"status": "Up",
"switch": "00:00:00:1a:1e:00:3b:40",
"version": "v1.3"
}
]
```

Host API

The Host API returns information about the hosts that are connected to the network, using the GET HTTP method.

Table 224: *Host API Query Parameters*

Query Parameter	Definition
host	Lists all hosts in the network.
host?mac	Lists all hosts associated with a specific MAC address.
host?ip	Lists all hosts associated with a specific IP address.
host?ip & timeout	Lists all hosts associated with a specific IP address and timeout period.
host?start	Controls paging of results by pointing to the starting object.

Query Parameter	Definition
host?limit	Controls paging of results by limiting the number of objects to be returned.
host?direction	Controls paging of results by setting the direction to *next* or *prev*.

The output for this message type displays the following information:

Table 225: *Host API Output Parameters*

Output Parameter	Definition
attachment-point	Information on the host's attachment point.
port	Number of the port to which the host is connected.
port-mac	MAC address of the port to which the host is connected.
switch	MAC address of the switch to which the host is connected.
created-at	Time that the host entry is created.
host-mac	MAC address of the host.
idle-for	Amount of time that host is idle, in seconds.
ip-addrs	IP address of the host.
up-time	Amount of time that the host is up, in seconds.
wireless	Indicates if the host is wireless.
object-count	Number of objects to be returned.
page-info	Information about result pages.
next-id	ID number of the next page.
previous-id	ID number of the previous page.
response-time	Query response time, in microseconds.

The Host API displays output similar to the example below:

```
#curl --insecure -b aruba-cookie https://10.4.251.105:4343/sdn/v1/host
[
{
"attachment-point": {
"port": 3,
"port-mac": "00:1a:1e:00:3b:46",
"switch": "00:00:00:1a:1e:00:3b:40"
},
"created-at": "Mon Jun 29 01:36:51 2015\n",
"host-mac": "00:00:5e:00:01:01",
"idle-for": 39464,
"ip-addrs": [
"63.82.214.201"
],
"up-time": 40037,
"wireless": false
```

```
}
]
"Meta-Info": {
"object-count": 4,
"page-info": {
"next-id": "573b0fd3cf42dcb5ab9c7c40"
"previous-id": "573b0fd3cf42dcb5ab9c7c20"
},
"reponse-time (micros)": 607
}
}
```

Flows API

The Flows API returns information about the flows between applications and data-plane elements, using the HTTP GET method.

Table 226: *Flows API Query Parameters*

Query Parameter	Definition
flows	Lists all flows between applications and data-plane elements.
flows?flow-group-id	Lists all flows associated with a specific flow-group.
flows?flow-id	Lists all flows associated with a specific flow ID.
flows?dpid	Lists all flows associated with a specific datapath ID.
flows?start	Controls paging of results by pointing to the starting object.
flows?limit	Controls paging of results by limiting the number of objects to be returned.
flows?direction	Controls paging of results by setting the direction to *next* or *prev*.

The output for this message type displays the following information:

Table 227: *Flows API Output Parameters*

Output Parameter	Definition
actions	Action(s) used by the flow. See Table 226 for the complete list of flow actions.
byte-count	Total byte count of the flow.
cookie	Cookie ID to which the flow is assigned.
created-at	Time that the flow entry is created.
dst-ip	Destination IP address of the flow.
ether-type	Ethertype used by the ether header.
flow-group-id	Flow-group ID.
flow-id	Flow ID.
hard-timeout	Hard-timeout, after which a flow is deleted based on the time period since the flow was created.

Output Parameter	Definition
idle-timeout	Idle-timeout, after which a flow is deleted based on the time period since activity was last detected.
packet-count	Number of packets transmitted by the flow.
priority	Priority of the flow.
protocol	Protocol used by the flow (for example, TCP).
src-ip	Source IP address of the flow.
status	Status of the flow (for example, install-confirmed).
switch	MAC address of the switch to which the flow is connected.
object-count	Number of objects to be returned.
page-info	Information about result pages.
next-id	ID number of the next page.
previous-id	ID number of the previous page.
response-time	Query response time, in microseconds.

The Flows API displays output similar to the example below:

```
#curl --insecure -b aruba-cookie https://10.4.251.105:4343/sdn/v1/flows
[
{
"actions": "output=controller",
"byte-count": 0,
"cookie": 281474976710660,
"created-at": "Mon Jun 29 07:54:10 2015\n",
"dst-ip": "2.2.2.2",
"ether-type": 2048,
"flow-group-id": 1007117466670727170,
"flow-id": 1007117466670791846,
"hard-timeout": 0,
"idle-timeout": 0,
"packet-count": 0,
"priority": 65535,
"protocol": 97,
"src-ip": "1.1.1.1",
"status": "Install-Confirmed",
"switch": "00:00:00:1a:1e:00:3b:40"
}
]
"Meta-Info": {
"object-count": 4,
"page-info": {
"next-id": "573b0fd3cf42dcb5ab9c7c40"
"previous-id”: “573b0fd3cf42dcb5ab9c7c20”
},
"reponse-time (micros)": 607
}
}
```

Links API

The Links API returns information about the inter-switch links that create the network topology, using the HTTP GET method.

Table 228: *Links API Query Parameters*

Query Parameter	Definition
links	Lists all links that make up the network topology.

The output for this message type displays the following information:

Table 229: *Links API Output Parameters*

Output Parameter	Definition
from-port	Port number of the source switch.
from-switch	MAC address of the source switch.
status	Status of the inter-switch link.
to-port	Port number of the destination switch.
to-switch	MAC address of the destination switch.

The Links API displays output similar to the example below:

```
#curl --insecure -b aruba-cookie https://10.4.251.201:4343/sdn/v1/links
[
{
"from-port": 2,
"from-switch": "00:00:00:00:00:00:00:03",
"status": 1,
"to-port": 1,
"to-switch": "00:00:00:00:00:00:00:04"
}
]
```

Path API

The Path API returns information on the path between two ports (hosts) in the network, using the GET HTTP method. Query the Host API to find the attachment point for each host, and then query the Path API to locate the path between the two attachment points.

Table 230: *Path API Query Parameters*

Query Parameter	Definition
path?src-spid&src-port&dst-dpid&dst-port	Lists a series of datapath IDs and ports, which constitute the path between the source datapath ID and port, and destination datapath ID and port.

The output for this message type displays the following information:

Table 231: *Path API Output Parameters*

Output Parameter	Definition
inPort	Ingress port of a switch in the path.
outPort	Egress port of a switch in the path.
switchDPID	Datapath ID of a switch in the path.

```
The Path API displays output similar to the example below:
#curl --insecure -b aruba-cookie
https://10.4.251.105:4343/sdn/v1/path?00:00:00:1a:1e:00:3b:40/1/00:00:00:1a:1e:00:3b:
90/2
[
{
"inPort": 1,
"outPort": 2,
"switchDPID": "00:00:00:1a:1e:00:3b:40"
},
{
"inPort": 3,
"outPort": 2,
"switchDPID": "00:00:00:1a:1e:00:3b:90"
}
]
```

Flows API

The Flows API installs flows between applications and data-plane elements, using the HTTP POST method. Refer to Flows API (GET HTTP) to view the list of parameters that can be specified to install a new flow.

The output for this message type displays the following information:

Table 232: *Flows API Output Parameters*

Output Parameter	Definition
flow-group-id	Flow-group in which the new flows are installed.
flows	List of new flows installed on the controller.
flow-id	ID of the new flow.
status	Status of flow installation.

The Flows API displays output similar to the example below:

```
#Install a single flow.
# curl --insecure -b "aruba-cookie" -d '{"flows": [{"switch":
"00:00:00:1a:1e:00:3b:40", "name":"sdn-1", "priority":32768, "ether-type":2048, "src-
ip":"20.20.20.4", "dst-ip":"20.20.20.5", "src-port":5000, "dst-port":8000,
"protocol":17, "actions":"output=controller,output=normal"}]}'
https://10.4.251.105:4343/sdn/v1/flows |python -mjson.tool
[
{
"Flow-Group-Id": 6269292156276441089,
"Flows": [
{
"Flow-Id": 6269292156276441860
```

```
}
],
"Status": "Install-In-Progress"
}
]
```

Flow Update API

The Flow Update API updates the list of actions that are installed on an existing flow, using the HTTP POST method. See Flow Match and Actions for more information about match fields and actions.

Table 233: *Flow Update API Query Parameters*

Query Parameter	Definition
flow-group-id	Updates all flows associated with a specific flow-group ID.
flow-id	(Optional) Updates all flows associated with a specific flow ID. If this field is specified, only the flows that match both the flow-group ID and flow ID are updated.

The output for this message type displays the following information:

Table 234: *Flow Update API Output Parameters*

Output Parameter	Definition
flow-group-id	ID of the flow-group that is being updated.
flow-id	ID of the flow that is being updated.
status	Status of the flow update.

The Flow Update API displays output similar to the example below:

```
#curl --insecure -b "aruba-cookie" -d '{"flow-group-id":6269292156276441089, "flow-
id":6269292156276441860, "actions":"output=controller"}'
https://10.4.251.105:4343/sdn/v1/flowupdate
{
"flow-group-id": 6269292156276441089,
"flow-id": 6269292156276441860,
"status": "Install-In-Progress"
}
```

Flow Delete API

The Flow Delete API deletes flows from the controller, using the HTTP POST method.

Table 235: *Flow Delete API Query Parameters*

Query Parameter	Definition
flow-group-id	Deletes all flows associated with a specific flow-group ID.
flow-group-id all	Deletes all flows in the network.

The output for this message type displays the following information:

Table 236: *Flow Delete API Output Parameters*

Output Parameter	Definition
status	Indicates if the flows for the given flow-group have been deleted.

The Flow Delete API displays output similar to the example below:

```
#curl --insecure -b "aruba-cookie" -d '{"flow-group-id":"6269292156276441090"}'
https://10.4.251.105:4343/sdn/v1/flowdelete  |python -mjson.tool
{
"Status": "Deleted"
}
```

Error Messages

Error messages are returned if any SDN REST API experiences the following errors:

Table 237: *Host API Output Parameters*

Error Message	Description
API-Timeout	The API times out.
Switch-Not-Reachable	The switch or data-plane element cannot be reached.
Switch-No-Reply	The switch or data-plane element does not reply to the request.
Db-Connect-Failed	The controller fails to connect to a database.
Out-of-Memory	The system runs out of memory.
Host-Not-Found	The host cannot be located.
Switch-Not-Found	The switch or data-plane element cannot be located.
Link-Not-Found	The link cannot be located.
Invalid-Input	The input is invalid.
Send-Failed	Northbound applications are unable to send data to internal applications.
Recv-Failed	Northbound applications are unable to receive data from internal applications.
Connect-Failed	Northbound applications are unable to connect to internal applications.
Bind-Failed	Northbound applications are unable to bind to the local address.
Socket-Failed	The socket connection fails.
Listen-Failed	Northbound applications are unable to listen for data from internal applications.
Accept-Failed	Northbound applications are unable to accept connections from internal applications.
Duplicate-Flow	The system encounters a duplicate flow.

Error Message	Description
Flow-Group-Not-Found	The flow-group cannot be located.
JSON-Parse-Error	The system experiences a JSON parsing error.
Flow-Conflict	The flow manager locates a conflicting entry in the flow database.
Switch-Internal-Error	The switch or data-plane element experiences an internal error.
Flow-Logical-Error	The system experiences an error in the flow logic.
Too-Many-Flows	Too many flows are being pushed in a single flow setup request.
Object-Id-Not-Found	The object ID cannot be located.
Object-Id-Invalid	The object ID is invalid.
Invalid-Group-Owner	The group owner is invalid.
Invalid-Action	The desired action is invalid.
System-Max-Flow-Limit-Reached	The system reaches the maximum flow limit.

Asynchronous APIs

Asynchronous APIs are implemented through the ZeroMQ, which is a TCP-based open-source library that offers publish and subscribe services for server-to-client communication. Northbound applications subscribe to the desired topics based on the type of information that is required by the client. The controller publishes this information as events or packets. All ZeroMQ events, except packet-out events, are published by the Northbound API and sent to the respective northbound applications. The packet-out events are published by the northbound applications and sent to the packetin-dispatcher on the controller. See Packet Handling for more information on packet processing.

The following ZeroMQ APIs are available on the SDN Controller:

Table 238: *SDN ZeroMQ APIs*

Event Type	Possible Values	Data Structure
Switch State Change(EVENT_SWITCH)	STATE_UP STATE_DOWN STATE_UPDATE	typedef struct { uint16_t event; uint16_t len; } event_header_t; typedef struct { uint64_t datapath_id; uint8_t state; uint8_t pad[7]; } switch_event_t;
Port State Change (EVENT_PORT)	PORT_LINK_UP PORT_LINK_DOWN	typedef struct { uint16_t event; uint16_t len; } event_header_t; typedef struct {

Event Type	Possible Values	Data Structure
		uint64_t datapath_id; uint32_t port_no; uint8_t port_mac[ETH_ADDR_LEN]; uint8_t reason; uint8_t state; uint8_t pad[MAX_PORT_NAME_LEN]; } port_event_t;
Packet In	N/A	typedef struct { uint64_t flow_id; uint64_t datapath_id; uint32_t port_no; uint16_t vlan_vid; uint16_t len; } zmq_pkt_in_t; Followed by packet data
Packet Out	N/A	typedef struct { uint64_t datapath_id; uint32_t port_no; uint16_t len; } zmq_pkt_out_t; Followed by packet data

User Event	USER_EVENT_ADD
	USER_EVENT_DELETE
	USER_EVENT_UPDATE
	USER_EVENT_IP_AGEOUT
		#define MAX_IP_ADDRS 4
		typedef struct {
		union {
		uint32_t ipv4_addr;
		struct in6_addr ipv6_addr;
		};
		bool is_ipv6;
		} ip_addr;
		
		typedef struct sdn_host_t_ {
		unsigned char mac[6];
		uint64_t dpid;
		uint32_t port;
		bool wireless;
		unsigned char port_mac[6];
		ip_addr addrs[MAX_IP_ADDRS];
		time_t updated_at; time_t
		created_at; } sdn_host_t;
		
		typedef struct user_event_ {
		uint8_t event_type;
		sdn_host_t host;
		} user_event_t;

Link Event	LINK_EVENT_ADD LINK_EVENT_DELETE LINK_EVENT_UPDATE	typedef struct sdn_link_status_t_ { uint64_t from_dpid; uint64_t to_dpid; uint32_t from_port; uint32_t to_port; uint8_t status;

Event Type	Possible Values	Data Structure
		} sdn_link_status_t; typedef struct link_event_ { uint8_t event_type; sdn_link_status_t link_status; } link_event_t;
Flow Event	FLOW_EVENT_DELETE FLOW_EVENT_ERROR FLOW_EVENT_ADD FLOW_EVENT_UPDATE	typedef struct flow_event_ { uint64_t flow_group_id; uint64_t flow_id; uint8_t event_type; } flow_event_t;

Northbound Authentication

To secure communication between users and APIs, the Northbound API supports basic authentication using HTTPS. During HTTPS authentication, the client is required to provide a username and a password for each HTTPS request. The server can only carry out the request after the user is authenticated.

OpenFlow agent runs on network devices such as switches, routers, wireless controllers and APs. This interacts with a centralized SDN Controller using the OpenFlow protocol. The OpenFlow agent translates OpenFlow commands into device specific actions.

The three main functions of the OpenFlow agent are:

1. Discover the Hosts—Help the SDN Controller to discover all the hosts (endpoints) attached to Mobility Master.
2. Discover the Network—Help the SDN Controller to learn about the Mobility Master's interface and its connectivity to other devices in the network.
3. Program the Network—Accept OpenFlow commands and take appropriate actions for those commands.
4. Provides Statistics—Provide visibility to SDN Controller to export flow or port statistics.

For OpenFlow to be functional in a network, you must enable SDN Controller on the Mobility Master and OpenFlow agent on the required Managed devices. By default, OpenFlow is enabled on Mobility Master as well as the managed devices.

Enabling SDN Controller on Mobility Master

The following procedure describes how to configure the SDN Controller:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In **All Profiles**, expand **Controller Profile**. Select **openflow-controller**.
3. In **openflow-controller**, select the **ofc-state** check box.



You can also configure an auxiliary channel port to reduce bandwidth consumption and latency on the main channel. To view the current status of the auxiliary channel, execute the **show openflow-controller** switches command. The default port is 6633. For more information on auxiliary channels, see Auxiliary Channel Driver on page 1042.

4. (Optional) To configure an auxiliary channel port, enter the listening port in the **ofc auxiliary-channel-port** field.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands enable SDN Controller on Mobility Master:

```
(host) [mm] (config) #openflow-controller
(host) [mm] (openflow-controller) #openflow-controller-enable
(host) [mm] (openflow-controller) #write memory
```



You can also configure an auxiliary channel port to reduce bandwidth consumption and latency on the main channel. To view the current status of the auxiliary channel, execute the **show openflow-controller** switches command. The default port is 6633. For more information on auxiliary channels, see Auxiliary Channel Driver on page 1042.

The following CLI command configure an auxiliary channel port on the SDN Controller:

```
(host) [mm] (openflow-controller) #auxiliary-channel-port <port-num>
```

Configuring OpenFlow Agent on Managed devices

To enable OpenFlow agent, you must perform the following tasks on the managed device:

1. Enable OpenFlow profile on the managed device.
   a. Configure the SDN Controller IP address and listening port.
   b. Bind the user VLAN.
2. Enable OpenFlow for the required user roles and Virtual APs.

The following sections describe the procedures to configure OpenFlow agent on a managed device:

Enabling OpenFlow and Binding User VLAN

The following procedure describes how to configure the OpenFlow profile:

1. In the **Managed Networks** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In **All Profiles**, expand **Controller Profile**. Select **Openflow-profile**.
3. In **Openflow-profile**, select the **State** check box.
4. (Optional) Select the **Auxiliary State** check box and enter the auxiliary port number in the **Auxiliary Channel Port** field to enable OpenFlow auxiliary channel port.



Ensure that the auxiliary channel port configured on the managed device matches with the one configured on Mobility Master. The default port is 6633.

5. In **controller-ip**, enter the Mobility Master IP address and port number.
6. In **bind-vlan**, enter the OpenFlow VLAN to the current list.
7. Click **Save**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure and enable the OpenFlow profile:

```
(host) [md] (config) #openflow-profile
(host) [md] (Openflow-profile) #openflow-enable
(host) [md] (Openflow-profile) #controller-ip <master-ip> <port>
```

The following CLI commands configure an auxiliary channel port:

```
(host) [md] (Openflow-profile) #openflow-auxiliary-enable
(host) [md] (Openflow-profile) #auxiliary-channel-port <port>
```



Ensure that the auxiliary channel port configured on the managed device matches with the one configured on Mobility Master. The default port is 6633.

The following CLI commands bind user VLANs:

```
(host) [md] (Openflow-profile) #bind-vlan <list of vlan ids separated by comma>
(host) [md] (Openflow-profile) #write memory
```

Enabling OpenFlow in User Role and Virtual AP

The following procedure describes how to enable OpenFlow in the user-role and virtual AP:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Roles & Policies**.
2. In the **Roles** tab, select an existing role.
3. In the **Roles > <custom-role>** section, click **Show Advanced View**.
4. Under **More**, expand **Network**.
5. Select the **Open flow** check box.
6. Click **Submit**.
7. Navigate to **Configuration > System > Profiles**.
8. In **All Profiles**, expand **Wireless LAN > Virtual AP**. Select the **default** profile. This procedure uses the *default* profile.
9. In **Virtual AP profile**, expand **Advanced**.
10. Select the **Openflow Enable** check box.
11. Click **Save**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**

The following CLI commands enable OpenFlow for a user role:

```
(host) [md] (config) #user-role <user-role>
(host) [md] (config-submode)#openflow-enable
(host) [md] (config-submode)#write memory
```

The following CLI commands enable OpenFlow for a VAP:

```
(host) [md] (config) #wlan virtual-ap <virtual-ap>
(host) [md] (Virtual AP profile "<virtual-ap>") #openflow-enable
(host) [md] (Virtual AP profile "<virtual-ap>") #write memory
```

Verifying OpenFlow Configuration on Managed Device

The following CLI command verifies the OpenFlow profile configuration on managed device:

```
(host) [md] #show openflow-profile

Openflow-profile "default"
--------------------------
Parameter                                       Value
---------                                       -----
State                                           Enabled
Openflow mode                                   passive
Openflow version                                v1.3
controller-ip                                   10.16.125.115:6633
VLAN ID or range(s) of VLAN IDs                 1,124,400,600
openflow tls                                    Disabled
```

```
certificate-file                                  none
key-file                                          none
ca-certificate-file                               none
```

Verifying OpenFlow Configuration on Mobility Master

The following CLI command verifies the OpenFlow profile configuration on Mobility Master.

```
(host) [mynode] #show openflow-controller

openflow-controller
-------------------
Parameter                       Value         Set
---------                       -----         ---
ofc state                       Enabled
ofc host-ageout-time            300
ofc mode                        passive
ofc tls                         Disabled
ofc certificate-file            none
ofc key-file                    none
ofc ca-certificate-file         none
ofc port                        6633
ofc topology-discovery          Enabled
ofc auxiliary-channel-port      6633
```

Viewing OpenFlow Information

The following show commands are used to view the OpenFlow related information:

- show openflow debug—Displays the OpenFlow debug information
- show openflow flows—Displays all the flows that are plumbed
- show openflow ports—Lists all the OpenFlow ports
- show openflow controller— Displays the OpenFlow Controller information
- show openflow capabilities —Displays the system capabilities
- show openflow flow-table— Displays the OpenFlow table
- show openflow statistics—Displays the OpenFlow statistics information
- show datapath openflow session/acl— Displays the session or ACL actions
- show datapath acl— Displays ACLs with OpenFlow index
- show ip access-list— Displays ACLs with action as OpenFlow

The Loadable Service Module feature provides an infrastructure that allows users to dynamically upgrade or downgrade individual service modules without requiring an entire system reboot. Services are delivered as individual service packages containing the version and instructions for loading and running the service. Loadable service module is introduced in ArubaOS 8.0.0.0.

This section includes the following topics:

- Service Modules
- Upgrading a Service Module
- Troubleshooting

Service Modules

The following service modules are Loadable Service Module-capable, and the default service packages are bundled with the ArubaOS image:

- AirGroup
- AppRF
- ARM
- AirMatch
- Northbound API
- Unified communications manager
- WebCC
- WLAN management system

Service Packages

Every service module has a corresponding service package, which can be downloaded from the Aruba support site and installed on Mobility Master.

Upgrading a Service Module

Service modules must be upgraded if there is a bug in the existing module or a newer version of the module has been released. Patches are posted to the Aruba Support site, where users can view and download packages to upgrade a service.



After an ArubaOS image upgrades or downgrades, the non-default service packages are deleted.

Downloading a Service Package

The following procedure describes how to upgrade a service module on Mobility Master:

1. Obtain the required service package from the Aruba Support site.
2. In the **Mobility Master**node-hierarchy, navigate to **Maintenance > Software Management > Service Module Packages** in the WebUI.
3. Click the **Add** button at the bottom of the **Service Module Packages** table to add a new service package.
4. Under **Add Package**, select the **Access method** used to fetch the package. Configure the settings described in Table 239.

Table 239: *Load New Package Configuration Parameters*

Parameter	Description
Access method	Select the protocol to send the service package from the image server to Mobility Master: ■ FTP ■ Local file ■ SCP ■ TFTP ■ USB
Host IP address	Enter the IP address of the image server where the service package resides.
Image file name	Enter the exact service package name as residing on the image server. **NOTE:** On selecting the **Local file** option from the **Access method** field, upload the service package from your local file explorer.
Destination file name	Enter the destination service package name. As a best practice, keep the image name same as destination file name.
Username	Enter the username of the image server. **NOTE:** This option is only available if you select the FTP or SCP protocol in the **Access method** field.
Password	Enter the password of the image server. **NOTE:** This option is only available if you select the FTP or SCP protocol in the **Access method** field.

5. Click **Submit** to validate the package.

Activating the Service Package

The following procedure describes how to activate the service package:

1. In the **Mobility Master**node hierarchy, navigate to **Maintenance > Software Management > Service Module Packages**.
2. Select the new package from the **Service Module Packages** table. The **Service Module Packages > [name]** window appears at the bottom of the workscreen.
3. Set the **Status** to **Active** to activate the new service package.
4. Click **Submit**.

Removing a Service Package

The following procedure describes how to remove a service package:

1. In the **Mobility Master**node hierarchy, navigate to **Maintenance > Software Management > Service Module Packages**.
2. Select a package from the **Service Module Packages** table. The **trash** icon for the highlighted package appears.
3. Click the **trash** icon. When the package delete window opens, click **Delete**.
4. Click **Submit**.

Downloading a Service Package

One of the following CLI commands downloads a service package through an FTP, SCP, or TFTP server:

```
(host) [mynode] #upgrade-pkg copy ftp: <ftphost> <username> <filename> flash:
<destfilename>
(host) [mynode] #upgrade-pkg copy scp: <scphost> <username> <filename> flash:
<destfilename>
(host) [mynode] #upgrade-pkg copy tftp: <tftphost> <filename> flash: <destfilename>
```

Upon download, the Loadable Service Module performs the following compatibility checks to determine if the package is compatible with the running version of ArubaOS. If validation is successful, the installation process can continue. If validation is unsuccessful, the package is removed, and an error message appears.

- **Platform Check**: Determines if the package must run on a specific platform.
- **Version Check**: Determines if the package version matches the version of ArubaOS running on the system.

Activating the Service Package

The following CLI command installs and activates the service package:

```
(host) [mynode] #upgrade-pkg activate <packagename>
```

The service is halted and upgraded with the new service package, during which time the service is unavailable to all users. After the new package is installed and activated, the service restarts.

Viewing Service Packages

The following CLI command displays the downloaded and active Loadable Service Module service packages on Mobility Master:

```
(host) [mynode] #show packages
```

Packages that are not active can be removed using the **upgrade-pkg remove <packagename>** command.

Removing a Service Package

The following CLI command deletes a service package from Mobility Master:

```
(host) [mynode] #upgrade-pkg remove <packagename>
```

Troubleshooting

Execute the **show packages upgrade-history** command to view package installation logs:

```
(host) [mynode] #show packages upgrade-history
```

```
May 04 21:50:29 Copying files to airgroup dir
May 04 21:50:29 Creating symbolic link to mdns binary
May 04 21:50:29 Package default_airgroup_pkg installation was successfully
May 04 21:50:29 Copying files to ucm dir
May 04 21:50:29 Creating symbolic link to ucm binary
May 04 21:50:29 Package default_ucm_pkg installation was successfully
May 04 21:50:30 Copying files to wms dir
May 04 21:50:30 Creating symbolic link to wms binary
May 04 21:50:30 Package default_wms_pkg installation was successfully
May 04 21:50:30 Copying files to arm_cm dir
May 04 21:50:30 Creating symbolic link to arm binary
May 04 21:50:30 Package default_arm_cm_pkg installation was successfully
May 04 21:50:30 Copying files to web_cc dir
May 04 21:50:30 Creating symbolic link to web_cc binary
May 04 21:50:30 Package default_web_cc_pkg installation was successfully
May 04 21:50:30 Copying files to nbapi_helper dir
May 04 21:50:30 Creating symbolic link to nbapi_helper binary
May 04 21:50:30 Package default_nbapi_helper_pkg installation was successfully
May 04 21:50:31 Copying files to airmatch dir
May 04 21:50:31 Copying airmatch binary
May 04 21:50:31 Package default_airmatch_pkg installation was successfully
May 04 21:50:31 Copying files to appRF dir
May 04 21:50:31 Creating symbolic link to appRF binary
```

Execute the **show packages supported** command to view the packages supported on Mobility Master:

```
(host) [mynode] #show packages supported
Packages Supported
------------------
Package Name  Version
------------  -------
airgroup      1
ucm           1
wms           1
arm_cm        1
web_cc        1
nbapi_helper  1
airmatch         1
appRF         1
```

This chapter outlines the steps required to configure voice and video services on the Mobility Master for Voice over IP (VoIP) devices, including Apple FaceTime, Alcatel-Lucent New Office Environment (NOE), Microsoft Lync, Skype for Business, or Teams, Cisco Jabber, Cisco Skinny Call Control Protocol (SCCP), Spectralink SVP, SIP, H.323, Vocera, and Wi-Fi Calling. As video and voice applications are more vulnerable to delay and jitter, the network infrastructure must be able to prioritize video and voice traffic over data traffic.

This chapter includes the following topics:

- Voice and Video License Requirements on page 1072
- Configuring Voice and Video on page 1072
- Working with QoS for Voice and Video on page 1083
- UCC on page 1089
- Understanding Extended Voice and Video Features on page 1124

Voice and Video License Requirements

The voice and video services require PEFNG licenses on the Mobility Master. For complete details on the required licenses, refer to the *Aruba Mobility Master Licensing Guide*.

Configuring Voice and Video

This section describes the steps required to set up and configure voice features on the Mobility Master:

1. Set up net services
2. Configure roles
3. Configure firewall settings for voice and video ALGs
4. Configure other parameters depending on the need and environment



Assigning voice traffic to the high priority queue is recommended when deploying voice over WLAN networks.

Voice ALG and Network Address Translation

Voice ALGs in Aruba Mobility Master do not support NAT. This is due to the way NAT functions and the way IP addresses are embedded in the signaling messages. In a typical customer deployment, a call server is deployed within an internal network which eliminates the need for NAT.

In short, voice ALGs should not be enabled when voice clients are behind a NAT.

Setting up Net Services

You can either use the default net services and ports or you can create or modify net services.

Using Default Net Services

The following table lists the default net services and their ports:

Table 240: *Default Voice Net Services and Ports*

Net Service Name	Protocol	Port	ALG
svc-h323-tcp	TCP	1720	H.323
svc-h323-udp	UDP	1718, 1719	H.323
svc-noe	UDP	32512	NOE
svc-noe-oxo	UDP	5000	NOE
svc-sccp	TCP	2000	SCCP
svc-sips	TCP	5061	SIPS
svc-sip-tcp	TCP	5060	SIP
svc-sip-udp	UDP	5060	SIP
svc-svp	119	0	SVP
svc-vocera	UDP	5002	VOCERA

Creating Custom Net Services

You can use CLI to create or modify net services.

```
(host) [mynode] (config) #netservice
[service name] [protocol] [port] [alg]
```

To create an svc-noe service on UDP port 32522, enter:

```
(host) [mynode] (config) #netservice svc-noe udp 32522 alg noe
```

Configuring User Roles

In the user-centric network, the user role of a wireless client determines its privileges and the type of traffic that it can send or receive in the wireless network. You can configure roles for clients that use mostly data traffic, such as laptops, and roles for clients that use mostly voice traffic, such as VoIP phones. Although there are different ways for a client to derive a user role, in most cases the clients using data traffic are assigned a role after they are authenticated through a method such as 802.1X, VPN, or captive portal. The user role for VoIP phones may also be derived from the OUI of their MAC addresses or the SSID to which they associate. Refer to Roles and Policies on page 491 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. Managed Device provides default user roles for all voice services. You can do one of the following:

- Using the Default User Role on page 1073
- Creating or Modifying Voice User Roles on page 1074
- Using the User-Derivation Rules on page 1076

Using the Default User Role

Managed Device 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:

- global-sacl
- apprf-voice-sacl
- ra-guard
- sip-acl
- noe-acl
- svp-acl
- vocera-acl
- skinny-acl
- h323-acl
- dhcp-acl
- tftp-acl
- dns-acl
- icmp-acl
- http-acl
- https-acl
- skype4b-acl
- jabber-acl
- wificalling-acl
- voip-applications-acl

For more details on the default voice role, enter the following command in the Mobility Master:

```
(host) [mynode] #show rights voice
```

Creating or Modifying Voice User Roles

You can create roles for Facetime, H.323, Jabber, NOE, SCCP, Skype for Business, SIP, SVP, Vocera, and Wi-Fi calling ALGs. The following procedure describes how to configure user roles for any of the ALGs:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Roles & Policies**.
2. In the **Policies** tab, click **+** to add a policy.
3. For **Policy Name**, enter a name.
4. For **Policy Type**, select **Session**.
5. Click **Submit**.
6. Select the newly added policy.
7. In **Policies > <custom-policy>**, click **+** to add a new rule.
8. Select **Access Control** option as the rule type.
9. Click **OK**.
10. Under **Roles**, configure the following settings:
    a. For **IP version**, select **IPv4**.
    b. For **Source**, select **any**.

c. For **Destination**, select **any**.

d. For **Service/app**, select service, then the correct voice or video ALG service. See Table 241 and Table 242 for service names for all ALGs:

Table 241: *Services for ALGs*

ALG	Service Name
NOE	▪ svc-noe ▪ sip-noe-oxo
SIP	▪ svc-sip-tcp ▪ svc-sip-udp
SIPS	svc-sips
SVP	svc-svp
VOCERA	svc-vocera
SCCP	svc-sccp
H.323	▪ svc-h323-tcp ▪ svc-h323-udp

Table 242: *Other Services for the ALGs*

ACL	Service Name
DHCP	svc-dhcp
TFTP	svc-tftp
ICMP	svc-icmp
DNS	svc-dns

e. For **Action**, select **permit**.

f. For **802.1p priority**, select a value. -- denotes lowest priority. 7 denotes highest priority.

g. Click **Submit**. Repeat steps 1 to 5 to add ACLs for more VoIP protocols.

11. Select the **Roles** tab. Click **+** to add a user role.
12. In the **New Role** window, for **Name**, enter a name for the user role.
13. Click **Submit**.
14. Select the newly added role.
15. In the **Roles > <custom-role>** section, click **Show Advanced View**. Configure the following settings:

    a. Under **Policies**, click **+**.

    b. In the **Add Policy** window, select the **Add an existing policy** option.

c. In the **Policy name** drop-down list, select the previously-configured policy name.

d. Click **Submit**.

e. Under **Policies**, click **+**.

f. In the **Add Policy** window, select the **Add an existing policy** option.

g. In the **Policy name** drop-down list, select **control**.

h. Click **Submit**.

16. Click **Pending Changes**.
17. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure user roles for ALGs:

```
(host) [md] (config) #ip access-list session <policy-name>
(host) ^[md] (config-submode) #any any <service-name> permit queue high
```

The following CLI commands map the policy name to the user role:

```
(host) [md] (config) #user-role <role-name>
(host) ^[md] (config-submode) #access-list session <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

Using the User-Derivation Rules

The user role can be derived from the attributes of the client 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 OUI of the MAC address of the client. The following procedure describes how derive a role based on SSID:



User-derivation rules are executed *before* the client is authenticated.

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Authentication > User Rules**.
2. In **User Rules Summary**, click **+**.
3. In the **Add New User Rule** window, enter a name for the user rule and click **Submit**.
4. In **User Rules Summary**, select the name of the user rule to configure the rule set.
5. In **Rules-set**, click **+** and configure the following settings:
   a. For **Set type**, select **Role** from the drop-down list.
   b. For **Rule type**, select **ESSID**.
   c. For **Condition**, select **equals**.
   d. For **Value**, enter the SSID used for the phones.
   e. For **Roles**, select the user role previously created.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Run the following commands to derive a role based on SSID:

```
(host) [md] (config) #aaa derivation-rules user <name of rule-set>
(host) ^[md] (config-submode) #set role condition essid equals <ssid-name> set-value <The value that the role/VLAN should be set to>
```

Deriving Role Based on MAC OUI

The following procedure describes how to derive a role based on MAC OUI:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Authentication > User Rules**.
2. In **User Rules Summary**, click **+**.
3. In the **Add New User Rule** window, enter a name for the user rule and click **Submit**.
4. In **User Rules Summary**, select the name of the user rule to configure the rule set.
5. In **Rules-set**, click **+** and configure the following settings:
   a. For **Set type**, select **Role** from the drop-down list.
   b. For **Rule type**, select **MAC Address**.
   c. For **Condition**, select **contains**.
   d. 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)
   e. For **Roles**, select the user role previously created.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Run the following commands to derive a role based on MAC OUI:

```
(host) [md] (config) #aaa derivation-rules user <name of rule-set>
(host) ^[md] (config-submode) #set role condition macaddr contains <xx:xx:xx:xx:xx:xx> set-value <The value that the role/VLAN should be set to>
```

Additional Video Configurations

You can configure ArubaOS to reliably and efficiently stream video traffic over WLAN. This new method allows you to stream video traffic reliably without much distortion. To ensure that video data is transmitted reliably, dynamic multicast optimization techniques are used.

Although the dynamic multicast optimization conversion generates more traffic, that traffic is buffered by the AP and delivered to the client when the client emerges from power-save mode.

The following sections describe:

- Configuring Video over WLAN enhancements on page 1077
- Prerequisites on page 1078

Configuring Video over WLAN enhancements

To configure video over WLAN enhancements:

- Enable **WMM** in the WLAN SSID profile.
- Enable **IGMP** proxy or IGMP snooping in the interface VLAN.
- Configure an ACL to set a DSCP value same as the **wmm-vi-dscp** value in the WLAN SSID profile for prioritizing the multicast video traffic.
- Enable **dynamic multicast optimization** in the virtual AP profile.

- Configure the **dynamic multicast optimization threshold** in the virtual AP profile. The maximum number of high throughput stations in a multicast group. The optimization will stop if the number exceeds the threshold value.
- Enable **multicast rate optimization** in the WLAN SSID profile to support higher data rate for multicast traffic in the absence of dynamic multicast optimization. Dynamic multicast optimization takes precedence over multicast rate optimization up to the configured threshold value.



Configuring the **Video Multicast Rate Optimization** parameter overrides the configuration of **BC/MC Rate Optimization** parameter for VI-tagged multicast traffic. Multicast traffic that is not VI-tagged behaves the same with BC/MC as before. If multicast rate is not set, all traffic behaves the same.

- Enable **video aware scan** in the ARM profile. This ensures that AP does not scan when a video stream is active.
- Optionally, you can configure and apply the **WMM bandwidth management profile** in the virtual AP profile. The total bandwidth share should not exceed 100 percent.
- Enable **multicast shaping** in the WMM bandwidth management profile to shape the sudden traffic from the source.

Prerequisites

You will need the Policy Enforcement Firewall Next Generation (PEFNG) license to enable dynamic multicast optimization. The following procedure describes how to configure video over WLAN enhancements:

1. Enable IGMP proxy or IGMP snooping on the managed device. To enable IGMP proxy:
   a. In the **Managed Network** node hierarchy, navigate to **Configuration > Interfaces > VLANs**.
   b. In **VLANs**, select an existing VLAN.
   c. In **VLANs > <vlan-name>**, select an existing VLAN ID.
   d. In the **IPv4** tab, expand **IGMP** and select **proxy** from the **IGMP** drop-down list.
   e. In **Proxy Interface**, select the **Interface** radio button and the appropriate interface from the drop-down list.
   f. Click **Submit**.
   g. Click **Pending Changes**.
   h. In the **Pending Changes** window, select the check box and click **Deploy changes**.

   To enable IGMP snooping:
   a. In the **Managed Network** node hierarchy, navigate to **Configuration > Interfaces > VLANs**.
   b. In **VLANs**, select an existing VLAN.
   c. In **VLANs > <vlan-name>**, select an existing VLAN ID.
   d. In the **IPv4** tab, expand **IGMP** and select **snooping** from the **IGMP** drop-down list.
   e. Click **Submit**.
   f. Click **Pending Changes**.
   g. In the **Pending Changes** window, select the check box and click **Deploy changes**.

2. Enable wireless multimedia and set a DSCP value for video traffic:
   a. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
   b. In **All Profiles**, expand **Wireless LAN > SSID**. Select the **default** profile.
      This example uses the *default* profile.
   c. In **SSID Profile**, select the **Wireless Multimedia (WMM)** check box.
   d. In the **DSCP mapping for WMM video AC (0-63)** field, enter the DSCP value (integer number).
   e. Click **Submit**.
   f. Click **Pending Changes**.
   g. In the **Pending Changes** window, select the check box and click **Deploy changes**.
3. Create an ACL on the managed device with the values equivalent to the DSCP mappings to prioritize the video traffic:
   a. In the **Managed Network** node hierarchy, navigate to **Configuration > Roles & Policies**.
   b. In the **Policies** tab, click **+** to add a policy.
   c. Enter the appropriate values in **Policies > <custom-policy>** to match the DSCP mapping values.
   d. Click **Pending Changes**.
   e. In the **Pending Changes** window, select the check box and click **Deploy changes**.

   You can also add this ACL to any user role or port. To apply the ACL to a user role:
   a. In the **Managed Network** node hierarchy, navigate to **Configuration > Roles & Policies**.
   b. Select the **Roles** tab and click **+** to add a user role.
   c. In the **New Role** window, for **Name**, enter a name for the user role.
   d. Click **Submit**.
   e. Select the newly added role.
   f. In the **Roles > <custom-role>** section, click **Show Advanced View**. Configure the following settings:
   g. Under **Policies**, click **+**.
   h. In the **Add Policy** window, select the **Add an existing policy** option.
   i. In the **Policy name** drop-down list, select the previously-configured policy name.
   j. Click **Submit**.
   k. Click **Pending Changes**.
   l. In the **Pending Changes** window, select the check box and click **Deploy changes**.

   To apply the ACL to a port:
   a. In the **Managed Network** node hierarchy, navigate to **Configuration > Interfaces> Ports**.
   b. In **Ports**, select an upstream port.
   c. Under the **VLAN Policy** drop-down list, select the ACL.

d. Click **Submit**.

e. Click **Pending Changes**.

f. In the **Pending Changes** window, select the check box and click **Deploy changes**.

4. Configure dynamic multicast optimization for video traffic on a virtual AP profile:

a. In the **Managed Network** node hierarchy, navigate to **Configuration > System> Profiles**.

b. In **All Profiles**, expand **Wireless LAN > Virtual AP**. Select the **default** profile.
This example uses the *default* profile.

c. In **Virtual AP profile**, expand **Broadcast/Multicast**.

d. Select the **Dynamic Multicast Optimization (DMO)** check box.

e. Click **Submit**.

f. Click **Pending Changes**.

g. In the **Pending Changes** window, select the check box and click **Deploy changes**.

5. Configure multicast rate optimization for the video traffic:

a. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.

b. In **All Profiles**, expand **Wireless LAN > SSID**. Select the **default** profile.
This example uses the *default* profile.

c. In **SSID Profile**, select the **BC/MC Rate Optimization** check box.

d. Select a value from the **Video Multicast Rate Optimization** drop-down list.

e. Click **Submit**.

f. Click **Pending Changes**.

g. In the **Pending Changes** window, select the check box and click **Deploy changes**.



Configuring the **Video Multicast Rate Optimization** parameter overrides the configuration of **BC/MC Rate Optimization** parameter for VI-tagged multicast traffic. Multicast traffic that is not VI-tagged behaves the same with BC/MC as before. If multicast rate is not set, all traffic behaves the same.

6. Configure ARM scanning for video traffic:

a. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.

b. In **All Profiles**, expand **RF Management > Adaptive Radio Management (ARM)**. Select the **default-a** profile.
This example uses the *default-a* profile.

c. In **Adaptive Radio Management (ARM) profile**, expand **Scanning** and select the **VoIP Aware Scan** check box.

d. Click **Submit**.

e. Click **Pending Changes**.

f. In the **Pending Changes** window, select the check box and click **Deploy changes**.

7. Configure and apply bandwidth management profile:
   a. In the **Managed Network** node hierarchy, navigate to **Configuration > System> Profiles**.
   b. In **All Profiles**, expand **QOS > WMM Traffic management**.
   c. In **WMM Traffic management profile: New Profile**, click **+** and enter a profile name.
   d. Select the **Enable Shaping Policy** check box, and enter the bandwidth share values across voice, video, best effort, and background.
   e. Click **Submit**.
   f. Click **Pending Changes**.
   g. In the **Pending Changes** window, select the check box and click **Deploy changes**.

   This step is optional.



Ensure that you configure the WMM traffic management profile to the virtual AP profile, if you have configured the virtual AP traffic management profile.

The bandwidth share percentage configuration of WMM traffic management feature is not supported on AP-203H, AP-203R, 207 Series, 200 Series, 210 Series, 220 Series, AP-228, 270 Series, 340 Series, 500 Series, 510 Series, 570 Series, and AP-518 access points.

   After you configure the WMM bandwidth management profile, apply it to the virtual AP profile.
8. Enable multicast shaping on the firewall:
   a. In the **Managed Network** node hierarchy, navigate to **Configuration > Services> Firewall**.
   b. In **Global Settings**, select the **Multicast automatic shaping** check box.
   c. Click **Submit**.
   d. Click **Pending Changes**.
   e. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the video over WLAN enhancements:

1. Enable IGMP proxy or IGMP snooping on the managed device.

   To enable IGMP proxy:

   ```
   (host) [md] (config) #interface vlan <id>
   (host) ^[md] (config-submode)#ip igmp proxy gigabitethernet <slot/module/port>
   ```

   To enable IGMP snooping:

   ```
   (host) [md] (config) #interface vlan <id>
   (host) ^[md] (config-submode)#ip igmp snooping
   ```
2. Enable wireless multimedia and set a DSCP value for video traffic:

   ```
   (host) [md] (config)#wlan ssid-profile default
   (host) ^[md] (SSID Profile "default")#wmm
   (host) ^[md] (SSID Profile "default")#wmm-vi-dscp <value>
   ```

   Setting the DSCP value tags the content as video stream that the APs can recognize.
3. Create an ACL on the managed device with the values equivalent to the DSCP mappings to prioritizes video traffic. The following ACL prioritizes the multicast traffic from the specified multicast group on the managed device. You can also add this ACL to any user role or port:

   ```
   (host) [md] (config) #ip access-list session mcast_video_acl
   ```

```
(host) ^[md] (config-submode)#any network 224.0.0.0 255.0.0.0 any permit tos 40
queue high dot1p-priority 5
```

a. To apply the ACL to a user role:

This example uses the user role *authenticated*.

```
(host) [md] (config) #user-role authenticated access-list session mcast_
video_acl
```

b. To apply the ACL to a port:

```
(host) [md] (config) #interface gigabitethernet <slot/module/port>
(host) ^[md] (config-submode)#ip access-group mcast_video_acl session
```

4. Configure dynamic multicast optimization for video traffic on a virtual AP profile:

```
(host) [md] (config)#wlan virtual-ap default
(host) ^[md] (Virtual AP Profile "default")#dynamic-mcast-optimization
```

5. Configure the dynamic multicast optimization threshold value:

```
(host) ^[md] (Virtual AP Profile "default")#dynamic-mcast-optimization-thresh 6
```

6. Configure multicast rate optimization for video traffic:

```
(host) [md] (config) #wlan ssid-profile default
(host) ^[md] (SSID Profile "default") #mcast-rate-opt
```

7. Configure ARM scanning for video traffic:

In the **rf arm-profile**, enable the **video-aware-scan** option. This prevents APs from scanning when a video traffic is active:

```
(host) [md] (config) #rf arm-profile default-a
(host) ^[md] (Adaptive Radio Management (ARM) profile "default-a") #video-aware-
scan
```

8. Configure and apply a bandwidth management profile:

Ensure that you configure the WMM traffic management profile to the virtual AP profile, if you have configured the virtual AP traffic management profile.



The bandwidth share percentage configuration of WMM traffic management feature is not supported on AP-203H, AP-203R, 207 Series, 200 Series, 210 Series, 220 Series, AP-228, 270 Series, 340 Series, 500 Series, 510 Series, 570 Series, and AP-518 access points.

a. Enable a bandwidth shaping policy so that the allocated bandwidth share is appropriately used:

```
(host) [md] (config) #wlan wmm-traffic-management-profile default
(host) ^[md] (WMM Traffic management profile "default") # enable-shaping
```

b. Set a bandwidth percentage for the following categories:

```
(host) ^[md] (WMM Traffic management profile "default") # background 10
(host) ^[md] (WMM Traffic management profile "default") # best-effort 20
(host) ^[md] (WMM Traffic management profile "default") # video 50
(host) ^[md] (WMM Traffic management profile "default") # voice 20
```

After you configure the WMM bandwidth management profile, apply it to the virtual AP profile:

```
(host) [md] (config) #wlan virtual-ap default
(host) ^[md] (Virtual AP profile "default") #wmm-traffic-management-profile
default
```

9. Enable multicast shaping on the firewall:

```
(host) [md] (config) #firewall
(host) ^[md] (config-submode) #shape-mcast
```

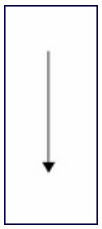
Working with QoS for Voice and Video

Quality of Service (QoS) settings for voice and video applications are configured when you configure firewall roles and policies.

Wi-Fi Multimedia (WMM), is a Wi-Fi Alliance specification based on the IEEE 802.11e wireless QoS standard. WMM works with 802.11a, b, g, n, and ac physical layer standards.

WMM supports four access categories (ACs); voice, video, best effort, and background. Table 243 shows the mapping of the WMM access categories to 802.1p priority values. The 802.1p priority value is contained in a two-byte QoS control field in the WMM data frame.

Table 243: *WMM Access Category to 802.1p Priority Mapping*

Priority	802.1p Priority	WMM Access Category
Lowest	1	Background
	2	
	0	Best effort
	3	
	4	Video
	5	
	6	Voice
Highest	7	

In non-WMM, or hybrid environments where some clients are not WMM-capable, ArubaOS 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 QoS data or Qos null data frame. For the environments in which the wireless clients support WMM, you can enable both WMM and U-APSD in the SSID profile.

Enabling WMM

The following procedure describes how to enable WMM for wireless clients:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System> Profiles**.
2. In **All Profiles**, expand **Wireless LAN > Virtual AP > default > SSID**.
   This example uses the *default* profile.
3. In **SSID Profile**, select the **Wireless Multimedia (WMM)** check box. Or, select the **Wireless Multimedia U-APSD (WMM-UAPSD) Powersave** check box if you want to enable WMM in power save mode.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure WMM for wireless clients:

```
(host) [md] (config) #wlan ssid-profile default
(host) ^[md] (SSID Profile "default") #wmm
(host) ^[md] (SSID Profile "default") #wmm-uapsd
```

You also configure WMM DSCP mapping in decrypt-tunnel.

Configuring WMM DSCP Mapping in Decrypt-tunnel

Starting from ArubaOS 8.4.0.0 WMM DSCP mapping supports IPv6 packets in the upstream direction of the decrypt tunnel mode. The WMM mapping for IPv4 and IPv6 is controlled by the DSCP mapping knob. When this knob is enabled in the decrypt-tunnel mode, the DSCP is mapped according to the WMM configuration for both IPv4 and IPv6 packets. When this knob is disabled in the decrypt-tunnel mode, there is no change to the DSCP mapping for both IPv4 and IPv6 packets. You can use the WebUI or CLI to configure DSCP mapping in decrypt-tunnel.

The following procedure describes how to enable the WMM DSCP Mapping:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System> Profiles.**
2. In **All Profiles**, expand **Wireless LAN > SSID**. Select the SSID profile.
3. In **SSID Profile** screen, select the **WMM DSCP Mapping Control** check box to enable this option. This feature is enabled by default.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command configures the WMM DSCP mapping:

```
(host) [mynode] #wmm-dscp-mapping
```

The following CLI command disables the WMM DSCP mapping:

```
(host) [mynode] #no wmm-dscp-mapping
```

Configuring WMM AC Mapping

The IEEE 802.11e standard defines the mapping between WMM ACs and Differentiated Services Code Point (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.

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 244 shows the default WMM AC to DSCP decimal mappings and the recommended WMM AC to DSCP mappings.

Table 244: *WMM Access Category to DSCP Mappings*

DSCP Decimal Value	WMM Access Category
8	Background
16	
0	Best effort
24	
32	Video
40	

DSCP Decimal Value	WMM Access Category
48	Voice
56	

By customizing WMM AC mappings, both the managed device 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 both upstream (client to AP) and downstream (AP to client) traffic.

When planning your mappings, make sure that any immediate switch or router does not have conflicting 802.1p or DSCP configuration or mapping. If this occurs, your traffic may not be prioritized correctly.

To view the mapping settings, run the following command:

```
(host) [mynode] #show wlan ssid-profile <profile>
```

The following procedure describes how to map WMM AC with DSCP:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System> Profiles**.
2. In **All Profiles**, expand **Wireless LAN > Virtual AP > default > SSID**.
   This example uses the *default* profile.
3. In **SSID Profile**, select the **Wireless Multimedia (WMM)** check box.
4. Modify the DSCP mapping settings, as needed:
   - **DSCP mapping for WMM voice AC (0-63)**—DSCP map for voice traffic
   - **DSCP mapping for WMM video AC (0-63)**—DSCP map for video traffic
   - **DSCP mapping for WMM best-effort AC (0-63)**—DSCP map for best-effort traffic
   - **DSCP mapping for WMM background AC (0-63)**—DSCP map for background traffic
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following commands map WMM AC with DSCP:

```
(host) [md] (config) #wlan ssid-profile <profile>
(host) ^[md] (SSID Profile "default") #wmm-be-dscp <wmm-be-dscp>
(host) ^[md] (SSID Profile "default") #wmm-bk-dscp <wmm-bk-dscp>
(host) ^[md] (SSID Profile "default") #wmm-vi-dscp <wmm-vi-dscp>
(host) ^[md] (SSID Profile "default") #wmm-vo-dscp <wmm-vo-dscp>
```

The WMM-DSCP mapping functionality has the following features:

- Default mappings are not there for a newly created SSID profile and for a factory default managed device running an ArubaOS 8.0 image.
- If the mapping has no value, the original DSCP for upstream traffic is retained.
- The maximum number of values that can be configured for WMM-DSCP is 8.
- For the upstream traffic, if the mapping exists and incoming DSCP value matches one of the mapped values, then the DSCP value is retained.
- For the upstream traffic, if the mapping exists and incoming DSCP value does not match any of the mapped values, then the DSCP value is overwritten with the first value in the WMM-DSCP list
- For wireless-to-wireless traffic, if the AC of the incoming packet has no mapping and the incoming DSCP value is mapped to a different AC, then the DSCP value is retained and WMM priority is changed to the corresponding AC where incoming DSCP is mapped.

Configuring DSCP Priorities

You can configure DSCP priorities for WMM packets in the following ways:

- Configure the DSCP mappings in the SSID profile
- Set a ToS value in the ACL
- Set the ToS value and the 802.1p priority in the ACL

Setting a ToS value in the ACL overrides the default DSCP mappings configured in the SSID profile. Configuring a DSCP priority in both the L2 and L3 header prioritizes the WMM packets with the higher value. For example, you can have different ToS values set for different voice traffic in a network. To prioritize all of them in the voice queue, we can set the 802.1p priority to voice.

Consider a deployment where Cisco Softphone, Microsoft Skype for Business, and Avaya Scopia are configured with the following DSCP:

- Cisco Softphone - DSCP 46
- Microsoft Skype for Business - DSCP 44
- Avaya Scopia - DSCP 42

In the absence of doing anything, all of the DSCP above would map into the video queue. To map all the traffic into voice queue, you can use the following ACL configuration:

```
(host) [md] (config) #wlan ssid-profile VOICE
(host) ^[md] (SSID Profile "VOICE") #wmm-vo-dscp 46
(host) ^[md] (SSID Profile "VOICE") #!
(host) ^[md] (config) #ip access-list session VOICE
(host) ^[md] (config-submode)#any destination <SKYPE4B_SERVER> <SKYPE4B_PORTS> permit tos
44 dot1p-priority 6
(host) ^[md] (config-submode)#any destination [SCOPIA_SERVER] [SCOPIA _PORTS] permit tos
42 dot1p-priority 6
```



You must know the ports on which each traffic is sent so that the correct traffic is identified.

Configuring Dynamic WMM Queue Management

Traditional wireless networks provide all clients with equal bandwidth access. However, delays or reductions in throughput can adversely affect voice and video applications, resulting in disrupted VoIP conversations or dropped frames in a streamed video. Thus, data streams that require strict latency and throughput need to be assigned higher traffic priority than other traffic types.

The Wi-Fi Alliance defined the Wi-Fi Multimedia (WMM) standard in response to industry requirements for QoS support for multimedia applications for wireless networks. This is defined as per the IEEE 802.11e standards.

WMM requires:

- the access point be Wi-Fi Certified and has WMM enabled
- the client device be Wi-Fi Certified
- the application support WMM

The following sections describe:

- Enhanced Distributed Channel Access on page 1087
- Configure EDCA Parameters on page 1087
- Configure EDCA Profile on page 1089

Enhanced Distributed Channel Access

WMM provides media access prioritization through Enhanced Distributed Channel Access (EDCA). EDCA defines four ACs to prioritize traffic; voice, video, best effort, and background. These ACs correspond to 802.1p priority tags, as shown in Table 245.

Table 245: *WMM Access Categories and 802.1p Tags*

WMM Access Category	Description	802.1p Tag
Voice	Highest priority	7, 6
Video	Prioritize video traffic above other data traffic	5, 4
Best Effort	Traffic from legacy devices or traffic from applications or devices that do not support QoS	0, 3
Background	Low priority traffic (file downloads, print jobs)	2, 1

While the WMM ACs designate specific types of traffic, you can determine the priority of the ACs. For example, you can choose to give video traffic the highest priority. With WMM, applications assign data packets to an AC. In the client, the data packets are then added to one of the transmit queues for voice, video, best effort, or background.

WMM is an extension to the Distributed Coordination Function (DCF) of the CSMA/CA protocol. The collision resolution algorithm responsible for traffic prioritization depends on the following configurable parameters for each AC:

- arbitrary inter-frame space number (AIFSN)
- minimum and maximum contention window (CW) size

For each AC, the backoff time is the sum of the AIFSN and a random value between 0 and the CW value. The AC with the lowest backoff time is granted the opportunity to transmit (TXOP). Frames with the highest-priority AC are more likely to get TXOP, because they tend to have the lowest backoff times (a result of having smaller AIFSN and CW parameter values). The value of the CW varies through time as the CW doubles after each collision up to the maximum CW. The CW is reset to the minimum value after successful transmission. In addition, you can configure the TXOP duration for each AC.

On the managed device, you configure the AC priorities in the WLAN EDCA parameters profile. There are two sets of EDCA profiles you can configure:

- AP parameters affecting traffic from the AP to the client.
- STA parameters affecting traffic from the client to the AP.

Configure EDCA Parameters

The following procedure describes how to define an Enhanced Distributed Channel Access (EDCA) profile for APs or for clients (stations):

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System> Profiles**.
2. In **All Profiles**, expand **Wireless LAN > SSID > default**.
   This example uses the *default* profile.
3. Select the **EDCA Parameters (AP)** or **EDCA Parameters (Station)** profile. Configure the EDCA profile based on the parameters described in Table 246.
4. Click **Submit**.

5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Table 246: *EDCA Parameters Station and EDCA Parameters AP Profile Settings*

Parameter	Description
Best Effort	Set the following parameters to define the best effort queue: ■ **aifsn**: arbitrary inter-frame space number. Range: 1-15. ■ **ecw-max**: the exponential (n) value of the maximum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 1-15. ■ **ecw-min**: the exponential (n) value of the minimum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 0-15. ■ **txop**: transmission opportunity, in units of 32 microseconds. Divide the desired transmission duration by 32 to determine the value to configure. For example, for a transmission duration of 3008 microseconds, enter 94 (3008/32). Range: 0-2047. ■ **acm**: this parameter specifies mandatory admission control. With a value of **1**, the client reserves the access category through traffic specification (TSPEC) signaling. A value of **0** disables this option.
Background	Set the following parameters to define the background queue: ■ **aifsn**: arbitrary inter-frame space number. Range: 1-15. ■ **ecw-max**: the exponential (n) value of the maximum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 1-15. ■ **ecw-min**: the exponential (n) value of the minimum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 0-15. ■ **txop**: transmission opportunity, in units of 32 microseconds. Divide the desired transmission duration by 32 to determine the value to configure. For example, for a transmission duration of 3008 microseconds, enter 94 (3008/32). Range: 0-2047. ■ **acm**: this parameter specifies mandatory admission control. With a value of **1**, the client reserves the access category through traffic specification (TSPEC) signaling. A value of **0** disables this option.
Video	Set the following parameters to define the background queue: ■ **aifsn**: trbitrary inter-frame space number. Range: 1-15. ■ **ecw-max**: The exponential (n) value of the maximum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 1-15. ■ **ecw-min**: the exponential (n) value of the minimum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 0-15. ■ **txop**: transmission opportunity, in units of 32 microseconds. Divide the desired transmission duration by 32 to determine the value to configure. For example, for a transmission duration of 3008 microseconds, enter 94 (3008/32). Range: 0-2047. ■ **acm**: this parameter specifies mandatory admission control. With a value of **1**, the client reserves the access category through traffic specification (TSPEC) signaling. A value of **0** disables this option.
Voice	Set the following parameters to define the background queue: ■ **aifsn**: trbitrary inter-frame space number. Range: 1-15.

Parameter	Description
	■ **ecw-max**: the exponential (n) value of the maximum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 1-15. ■ **ecw-min**: the exponential (n) value of the minimum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 0-15. ■ **txop**: transmission opportunity, in units of 32 microseconds. Divide the desired transmission duration by 32 to determine the value to configure. For example, for a transmission duration of 3008 microseconds, enter 94 (3008/32). Range: 0-2047. ■ **acm**: this parameter specifies mandatory admission control. With a value of **1**, the client reserves the access category through traffic specification (TSPEC) signaling. A value of **0** disables this option.

Configure EDCA Profile

The following CLI commands define an Enhanced Distributed Channel Access (EDCA) profile for APs or for clients (stations):

```
(host) [md] (config) #wlan edca-parameters-profile {ap|station} <profile>
(host) ^[md] (EDCA Parameters profile (AP) "default") #{background|best-
effort|video|voice} [acm][aifsn <number>] [ecw-max <exponent> [ecw-min <exponent>] [txop
<number>]
```

The following CLI commands help to associate the EDCA profile instance to a SSID profile:

```
(host) [md] (config) #wlan ssid-profile <profile>
(host) ^[md] (SSID Profile "<profile>") #edca-parameters-profile {ap|sta} <profile>
```

UCC

This section describes the UCC feature. The Unified Communications Manager (UCM) is the core solution component of this feature. UCC addresses the onslaught of mobile devices that use voice, video, and collaboration applications. UCC solution reduces the cost of infrastructure for enterprise communication and collaboration.

This section includes the following topics:

- UCC Application in ArubaOS on page 1090
- UCC Value Additions in Mobility Master on page 1090
- UCC Changes in ArubaOS 8.0.0.0 on page 1090
- UCC Features Deprecated in ArubaOS on page 1091
- Prerequisites to Enable UCC on page 1091
- Multi-ALG Support on page 1094
- UCC ALG Configuration on page 1094
- View UCC Information on page 1096
- Custom SIP on page 1097
- Intelligent Call Handling on page 1097
- AppRF Integration with ALGs and User Role on page 1100
- Microsoft Lync or Skype for Business on page 1102
- Cisco Jabber on page 1113
- Wi-Fi Calling on page 1118
- UCC Dashboard on page 1122

- UCC-AirWave Integration on page 1122
- UCC Limitations on page 1124

UCC Application in ArubaOS

Starting from ArubaOS 8.x, UCM runs as a loadable service module on Mobility Master. UCC supports various applications like Apple FaceTime, Alcatel-Lucent NOE, Microsoft® Lync or Skype for Business, Cisco Jabber, Cisco skinny call control protocol, SpectraLink voice priority, SIP, H.323, Vocera, and Wi-Fi Calling. UCC application on Mobility Master implements the VoIP ALG to support both encrypted and non-encrypted VoIP protocols. UCC application uses the OpenFlow infrastructure to receive the signaling messages from the managed devices and also install and delete flows on the managed devices for calls made.

In addition, UCC is supported on stand-alone controllers.

UCC Value Additions in Mobility Master

The following is a list of UCC value additions in Mobility Master:

- Enables VoIP ALGs to run as a service on Mobility Master and managed devices need not run the same. This results in better scalability.
- Enables real-time analysis of VoIP calls in upstream direction. This is the real-time analysis and UCC call quality statistics calculated based on VoIP stream captured at the managed device.
- Supports Loadable Service Module. UCM is a Loadable Service Module. ALGs are completely decoupled from the managed devices. This enables faster innovation of VoIP services such as introduction of new ALGs and enhancements to existing features as they will become independent of the ArubaOS release version.
- Provides a solution to the fanout problem in Lync or Skype for Business SDN API. In earlier ArubaOS versions, Lync or Skype for Business SDN Manager sent call information messages to every local controller in the network, regardless of whether the local controller is involved in the call or not. This additional processing is an unnecessary overhead on the local controller. In addition, the bandwidth utilization between the data center and remote location is not efficient. With the Mobility Master deployment, Lync or Skype for Business SDN Manager sends the call information messages to Mobility Master only.
- Provides aggregation of statistical information of call-related data at a centralized entity.

UCC Changes in ArubaOS 8.0.0.0

The following is a list of UCC changes ArubaOS 8.0.0.0:

- In earlier ArubaOS versions, VoIP ALGs run on the respective local controllers that parse the signaling messages, dynamically opens sessions in firewall, prioritizes traffic, and provide visibility. In ArubaOS 8.0.0.0, VoIP ALGs do not run on the managed devices. They run as an application on Mobility Master. In a stand-alone controller deployment, the VoIP ALGs run on the stand-alone controller itself.
- UCC running on Mobility Master uses OpenFlow infrastructure to receive signaling packets on Mobility Master, parse, open sessions in the firewall, and prioritize them.
- Visibility for all supported UCC applications are provided from the centralized Mobility Master dashboard. You should login to individual managed device to view dashboard information.
- ArubaOS 8.x supports Cisco Jabber and Wi-Fi Calling.
- Unlike earlier ArubaOS versions.X, where ALGs use WMM-DSCP mappings in the WLAN SSID profile to set the ToS for RTP and RTCP, Mobility Master has ALG-specific QoS configurations.

UCC Features Deprecated in ArubaOS

The following are the features deprecated in ArubaOS:

- BSS transition and force BSS transition.
- Call count, bandwidth, and TSPEC-based call admission control.
- Classify media action in ACL for media classification – Microsoft® Lync or Skype for Business calls will automatically get prioritized without the need for classify media ACLs.
- SIP session timer.
- SIP dial plans.
- WMM-DSCP override setting in the SSID profile.
- Stateful ALG settings in global firewall options. These settings are now available in Mobility Master and master controller under the **Configuration > System > Profiles > All Profiles > UCC** profile.
- Lync or Skype for Business traffic control profile.
- Web Server port configuration for Lync or Skype for Business SDN API. The Mobility Master and master controller uses 32000 as the default port now.
- The **Monitoring** tab in the WebUI.
- The **show voice** commands.
- **sip-authentication-role** parameter in AAA profile.
- **voice-aware** parameter in AAA authentication 802.1X profile.

Prerequisites to Enable UCC

This section describes the prerequisites to enable UCC.

- OpenFlow Configuration on page 1091
- OpenFlow Profile Configuration on Managed Devices on page 1091
- OpenFlow in User Role and Virtual AP Configuration on page 1092
- Management Server Profile Configuration on page 1092
- Deep Packet Inspection Configuration on page 1093
- Firewall Visibility Configuration on page 1094

OpenFlow Configuration

Enable OpenFlow on Mobility Master. You must enable this in the **/mm** node hierarchy.

The following procedure describes how to configure OpenFlow on Mobility Master:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In **All Profiles**, expand **Controller Profile**. Select **openflow-controller**.
3. In **openflow-controller**, select the **ofc-state** check box.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure OpenFlow on Mobility Master:

```
(host) [mm] (config) #openflow-controller
(host) ^[mm] (openflow-controller) #openflow-controller-enable
(host) ^[mm] (openflow-controller) #write memory
```

OpenFlow Profile Configuration on Managed Devices

Bind the user VLANs to the OpenFlow profile on the managed devices. You must bind this in the **/md** node hierarchy.

The following procedure describes how to bind the user VLANs to the OpenFlow profile on the managed devices:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In **All Profiles**, expand **Controller Profile**. Select **Openflow-profile**.
3. In **Openflow-profile**, select the **State** check box.
4. In **controller-ip**, enter the Mobility Master IP address and port number.
5. In **bind-vlan**, enter the user VLAN to the current list.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands bind the user VLANs to the OpenFlow profile on the managed devices:

```
(host) [md](config) #openflow-profile
(host) ^[md](Openflow-profile) #openflow-enable
(host) ^[md](Openflow-profile) #controller-ip <MM-ip> <port>
(host) ^[md](Openflow-profile) #bind-vlan <list of user vlans>
(host) ^[md](Openflow-profile) #write memory
```

OpenFlow in User Role and Virtual AP Configuration

Enable OpenFlow in the user-role and the virtual AP profile. You must enable this in the **/md** node hierarchy.

The following procedure describes how to enable OpenFlow in the user-role and virtual AP:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Roles & Policies**.
2. In the **Roles** tab, select an existing role.
3. In the **Roles > <custom-role>** section, click **Show Advanced View**.
4. Under **More**, expand **Network**.
5. Select the **Open flow** check box.
6. Click **Submit**.
7. Navigate to **Configuration > System > Profiles**.
8. In **All Profiles**, expand **Wireless LAN > Virtual AP**. Select the **default** profile.
   This example uses the *default* profile.
9. In **Virtual AP profile**, expand **Advanced**.
10. Select the **Openflow Enable** check box.
11. Click **Submit**.
12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands enable OpenFlow in the user-role and virtual AP:

```
(host) [md] (config) #user-role <user-role>
(host) ^[md] (config-submode)#openflow-enable
(host) ^[md] (config-submode)#!
(host) ^[md] (config) #wlan virtual-ap <virtual-ap>
(host) ^[md] (Virtual AP profile "<virtual-ap>") #openflow-enable
(host) ^[md] (Virtual AP profile "<virtual-ap>") #write memory
```

Management Server Profile Configuration

Configure the management server profile. This enables AMON feeds to be sent to Mobility Master or master controller for various statistics. You must configure this in the **/md** node hierarchy or the sub-nodes of **/md**.

The following procedure describes how to configure management server profile:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In **All Profiles**, expand **Controller Profile**. Select **Mgmt Config**.
3. In **Mgmt Config profile**, click **+**.
4. In the **Profile name** field, enter the name of the management server profile.
5. Select the following check boxes:
   a. **Stats**
   b. **Sessions**
   c. **Monitored Info - Add/Update**
   d. **Monitored Info - Deletion**
   e. **Monitored Info - Periodic Snapshot**
6. Click **Submit**.
7. Navigate to **Configuration > System > More**.
8. Expand **General** and in **MON Receivers**, click **+**.
9. In **New MON Receiver**, enter the following details:
   a. In the **Server** field, enter the Mobility Master or master controller IP address.
   b. In the **Profile list** drop-down list, select the newly created management server profile.
10. Click **Submit**.
11. Click **Pending Changes**.
12. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure management server profile:

```
(host) [md] (config) #mgmt-server profile <profile-name>
(host) ^[md] (Mgmt Config profile "<profile-name>") #stats-enable
(host) ^[md] (Mgmt Config profile "<profile-name>") #sessions-enable
(host) ^[md] (Mgmt Config profile "<profile-name>") #monitored-info-enable
(host) ^[md] (Mgmt Config profile "<profile-name>") #monitored-info-del-enable
(host) ^[md] (Mgmt Config profile "<profile-name>") #monitored-info-snapshot-enable
(host) ^[md] (Mgmt Config profile "<profile-name>") #!
(host) ^[md] (config) #mgmt-server primary-server <MM-IP> profile <profile-name>
(host) ^[md] (config) #write memory
```

Deep Packet Inspection Configuration

Enable DPI on the managed devices if your deployment has Cisco Jabber clients. You must enable this in the **/md** node hierarchy. The following procedure describes how to enable DPI on the managed devices:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Services > Firewall**.
2. Expand **Global Settings**.
3. Select the **Enable deep packet inspection** check box.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command enables DPI on the managed devices:

```
(host) [md] (config) #firewall
(host) ^[md] (config-submode)#dpi
(host) ^[md] (config) #write memory
```



If DPI is enabled, either there should be an explicit ACL to permit RTP or RTCP traffic or an app-based ACL to permit media traffic. For more information, see Deep Packet Inspection Configuration on page 1093.

Firewall Visibility Configuration

Enable firewall visibility on the managed devices. This is an optional setting. Enable this setting to view traffic analysis on the Mobility Master dashboard. You must enable this in the **/md** node hierarchy.



Firewall sessions for Lync or Skype for Business desktop-sharing and file-transfer are not allowed. Manually open a range of TCP ports under the user role to allow Lync or Skype for Business desktop-sharing and file-transfer traffic. To allow a specific range of ports in the user role, refer the Microsoft Technet article which describes the port ranges used by Lync or Skype for Business clients and servers. Before media transmission, a Lync or Skype for Business client initiates a Session Traversal Utilities for NAT (STUN) connectivity check. Sessions created by STUN are subjected to media classification that classifies the media as Real-time Transport Protocol (RTP) or non-RTP. The firewall automatically allows the RTP session on the managed device and denies the non-RTP sessions.

The STUN connectivity check is always enabled irrespective of whether firewall visibility or DPI is disabled.

The following procedure describes how to enable firewall visibility on the managed devices:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Services > Firewall**.
2. Expand **Global Settings**.
3. Select the **Enable firewall visibility** check box.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure firewall visibility on the managed devices:

```
(host) [md] (config) #firewall-visibility
(host) ^[md] (config) #write memory
```

Multi-ALG Support

In ArubaOS 8.x, multiple applications running simultaneously on the same client device can be identified and prioritized. A maximum of 10 applications running simultaneously on client device is supported. The multi-ALG feature is enabled by default on Mobility Master.

UCC ALG Configuration

The UCC ALGs must be configured from the **/mm** node hierarchy of Mobility Master. All the ALGs are enabled by default.



SpectraLink voice priority ALG is enabled by default. ArubaOS does not have a separate configuration setting for this ALG.

The following procedure describes how to configure the ALGs:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles > All Profiles**.
2. Expand **UCC** to configure various ALGs as described in Table 247.
3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Table 247: *ALG Configurations*

ALG	Description
FaceTime ALG Configuration	Configures the Apple FaceTime ALG. The ALG is enabled by default. The range is 0-63. The DSCP value for the video session is 34 by default.
H323 ALG Configuration	Configures the H.323 ALG. The ALG is enabled by default. The range is 0-63. The DSCP value for the voice session is 46 by default.
Intelligent Call Handling Configuration	Configures the Intelligent Call Handling. The setting is enabled by default. The range is 50-95. The Channel Utilization Threshold is 90 by default.
Jabber ALG Configuration	Configures the Cisco Jabber ALG. The ALG is enabled by default. Enter the Cisco Unified Communication Manager IM & Presence server IP. The range is 0-63. The DSCP values for the voice, video, and app-sharing sessions are 46, 34, and 34, respectively, by default.
NOE ALG Configuration	Configures the Alcatel-Lucent NOE ALG. The ALG is enabled by default. The range is 0-63. The DSCP value for the voice session is 46 by default.
Real-Time Analysis Configuration	Configures the real-time analysis of VoIP calls including upstream real-time analysis. The setting is enabled by default.
SCCP ALG Configuration	Configures the Cisco SCCP ALG. The ALG is enabled by default. The range is 0-63. The DSCP value for the voice session is 46 by default.
SIP ALG Configuration	Configures the SIP ALG. The ALG is enabled by default. You can enable the **SIP Midcall request timeout** and **RTCP inactivity** settings. The range is 0-63. The DSCP values for the voice and video sessions are 46 and 34, respectively, by default.
Skybe4B ALG Configuration	Configures the Microsoft® Lync or Skype for Business ALG. The ALG is enabled by default. You can set the Lync or Skype for Business SDN listen protocol over HTTP or HTTPS. The default Lync or Skype for Business SDN API listen port is 32000. Based on the SDN listen protocol configuration, Mobility Master accepts either HTTP or HTTPS messages from the Lync or Skype for Business SDN Manager. The DSCP values for the voice, video, and app-sharing sessions are 46, 34, and 34, respectively, by default. The range is 0-63.
UCC Session Idle Timeout Configuration	Configures the UCC session idle timeout. On configuring this parameter, if the voice session is idle for the configured period, UCM aborts the session on the managed device due to inactivity. The range is 35-250. The default value is 35.
Vocera ALG Configuration	Configure the Vocera ALG. The ALG is enabled by default. The range is 0-63. The DSCP value for the voice session is 46 by default.
Wi-Fi Calling Configuration	Configures the Wi-Fi Calling. Wi-Fi Calling is enabled by default. The range is 0-62. The DSCP value for the voice session is 46 by default.

Table 247: *ALG Configurations*

ALG	Description
	dns-pattern–Configure the DNS pattern for the carrier. A maximum of 10 DNS patterns can be configured. DNS patterns for known carriers are configured by default. Default built-in patterns are: ■ 3 HK - wlan.three.com.hk ■ ATT - epdg.epc.att.net ■ Rogers - epdg.epc.mnc720.mcc302.pub.3gppnetwork.org ■ SmarTone - epdg.epc.mnc006.mcc454.pub.3gppnetwork.org ■ Sprint - primgw.vowifi2.spcsdns.net ■ T-Mobile - ss.epdg.epc.mnc260.mcc310.pub.3gppnetwork.org ■ Verizon - wo.vzwwo.com If the ePDG FQDN of the carrier does not match with the default patterns, use this option to configure the DNS pattern for the carrier. **NOTE:** The DNS IP address that Mobility Master learns for Wi-Fi Calling age out automatically, if there was no DNS query or response matching that IP for more than seven days. **service-provider**–Add the service provider name for enhanced visibility.

The following CLI commands configure the ALGs:

```
(host) [mm] (config) #ucc ?
facetime                Configure the FaceTime ALG Configuration
h323                    Configure the H323 ALG Configuration
ich                     Configure the Intelligent Call Handling Configuration
jabber                  Configure the Jabber ALG Configuration
noe                     Configure the NOE ALG Configuration
rtpa-config             Configure the Real-Time Analysis Configuration
sccp                    Configure the SCCP ALG Configuration
session-idle-timeout    Configure the UCC Session Idle Timeout Configuration
sip                     Configure the SIP ALG Configuration
skype4b                 Configure the Skype4B ALG Configuration
vocera                  Configure the Vocera ALG Configuration
wificalling             Configure the WiFiCalling Configuration
```

For more information, see the **ucc** command in the *ArubaOS Command-Line Interface Reference Guide*.

View UCC Information

The following commands are available to view UCC client and call information using the CLI:

```
(host) [mm] #show ucc ?
call-info               Show ucc call detailed records (CDRs)
client-info             Show ucc client status and record
dns-ip-learning         DNS ip learning
facetime                Show the FaceTime ALG Configuration
h323                    Show the H323 ALG Configuration
ich                     Show the Intelligent Call Handling Configuration
internal-state          UCC internal-state information
jabber                  Show the Jabber ALG Configuration
noe                     Show the NOE ALG Configuration
rtpa-config             Show the Real-Time Analysis Configuration
rtpa-report             Show Real-Time Analysis report
sccp                    Show the SCCP ALG Configuration
session-idle-timeout    Show the UCC Session Idle Timeout Configuration
sip                     Show the SIP ALG Configuration
```

```
skype4b                 Show the Skype4B ALG Configuration
statistics              UCC statistics
trace-buffer            Show call trace buffer
vocera                  Show the Vocera ALG Configuration
wificalling             Show the WiFiCalling Configuration
```

For more information, see the **show ucc** commands in the *ArubaOS Command-Line Interface Reference Guide*.

Custom SIP

The UCC SIP application level gateway classifies and prioritizes the SIP media traffic only if the standard SIP port (5060) is used for SIP signaling. The custom SIP feature classifies and prioritizes the SIP media traffic that is compliant with SIP protocol but uses non-standard port for SIP signaling. The custom SIP feature uses a unique signature that is present in the SIP invite messages to classify and prioritize the SIP media traffic.

If you choose to use a non-standard port for SIP signaling, then configure the non-standard port by using the **netservice** command. This ensures that the non-standard port is classified for SIP signaling. If you want to display the specific custom-SIP application name in the CDRs and provision the SIP media for explicit voice or video priority exclusively, then use the custom SIP feature. The custom SIP feature can be enabled only if SIP is enabled and the non-standard port is configured for SIP ALG through the **netservice** commands on the managed device.

The following command configures the custom SIP ALG using the CLI:



Configure the custom SIP ALG from the **/mm** node hierarchy of Mobility Master.

```
(host) [mm] (config) #ucc custom-sip
app-name                app name
custom_sip_port         UCC custom-sip port
enable                  Enable Custom SIP ALG
key                     User-Agent:Key-name
no                      Delete command
priority                UCC Session Priority Config
```

The following command configures the netservice using the CLI:

```
(host) [mm] (config) #netservice test tcp 55060 alg sip
```

The following command shows the custom SIP ALG configuration using the CLI:

```
(host) [mm] (config) #show ucc custom-sip

custom-sip ALG Configuration
----------------------------
Parameter        Value
---------        -----
SIP ALG Support  Enabled
app-name         test
key-name         ConnectVoice
voice priority   34
video priority   48
custom-sip-port  55060
```

Intelligent Call Handling

ArubaOS 8.x replaces Call Admission Control with Intelligent Call Handling (ICH). ICH monitors the channel utilization of all radios of the APs on the managed device. If the channel utilization exceeds beyond a configurable threshold on a radio, new UCC calls are not prioritized. This is to ensure that existing calls on the radio are not penalized due to a new call when channel utilization is very high. ICH is enabled by default and applies to all ALGs supported by UCM.

The following procedure describes how to configure ICH:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles > All Profiles**.
2. Expand **UCC** and click **Intelligent Call Handling Configuration**.
3. In the **Intelligent Call Handling Configuration** section, configure the following settings:
4. Select the **Intelligent Call Handling** check box.
5. In the **Channel Utilization Threshold** text-box, enter the channel utilization value.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure ICH:

```
(host) [mm] (config) #ucc ich
(host) ^[mm] (Intelligent Call Handling Configuration) #enable
(host) ^[mm] (Intelligent Call Handling Configuration) #channel-utilization-threshold 90
```

RTP Analysis

Mobility Master performs RTP analysis for most VoIP ALG calls in both downstream (at AP) and upstream direction (at managed device) and captures the quality metrics. The downstream UCC score measures call quality between the AP and the wireless client in the downstream direction. The upstream UCC score measures call quality over the wired network between the AP and the managed device in the upstream direction. The quality metrics captured is applicable for all the active sessions belonging to the same or different ALG running on that client.

Starting with ArubaOS 8.1.0.0, Mobility Master calculates upstream UCC score for wired clients that are behind the wired port of an AP or Remote AP.

The following procedure describes how to configure RTP analysis:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles > All Profiles**.
2. Expand **UCC** and click **Real-Time Analysis Configuration**.
3. In the **Real-Time Analysis Configuration** section, configure the settings described in Real-Time Analysis Configuration Parameters.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Table 248: *Real-Time Analysis Configuration Parameters*

Parameter	Description
Real-Time Analysis of VoIP calls	Enables real-time analysis of VoIP calls. This is the real-time analysis and UCC statistics calculated based on VoIP stream at the access point.
Upstream Real-Time Analysis of VoIP calls	Enables real-time analysis of upstream VoIP calls. This is the real-time analysis and UCC statistics calculated based on VoIP stream at the managed device.

The following CLI commands configure real-time analysis:

```
(host) [mm] (config) #ucc rtpa-config
(host) ^[mm] (Real-Time Analysis Configuration) #enable
```

```
(host) ^[mm] (Real-Time Analysis Configuration) #upstream
(host) ^[mm] (Real-Time Analysis Configuration) #write memory
```



The upstream and downstream RTP analysis of VoIP calls are enabled by default.

The following CLI command displays the real-time analysis configuration:

```
(host) [mm] #show ucc rtpa-config

Real-Time Analysis Configuration
--------------------------------
Parameter                                  Value    Set
---------                                  -----    ---
Real-Time Analysis of VoIP calls           Enabled
Upstream Real-Time Analysis of VoIP calls  Enabled
```

The following CLI command displays the real-time analysis report:

```
(host) [mm] #show ucc rtpa-report

Help: [C] - Metric calculated at the Controller
      [A] - Metric calculated at the AP
      [E] - Metric calculated End-to-End

Real-Time Analysis Call Quality Report
--------------------------------------
Client(IP)       Client(MAC)        Client(Name)  ALG  Jitter(usec)[C]  Pkt-loss(%)[C]
Delay(usec)[C]
----------       -----------        ------------  ---  ---------------  --------------  -
-------------
192.168.201.240  f0:7b:cb:3b:65:5c  1002          SIP  23.700           0.000
101.800
192.168.201.246  00:24:d7:40:a8:58  1003          SIP  30.912           0.000
257.140

UCC Score[C]  Jitter(usec)[A]  Pkt-loss(%)[A]  Delay(usec)[A]  UCC Score[A]  Forward mode
------------  ---------------  --------------  --------------  ------------  ------------
68.366        0.000            0.499           316.400         84.119        decrypt-
tunnel
82.551        0.000            0.000           327.478         85.999        decrypt-
tunnel

Num Records:2
```

The following command displays real-time analysis for VoIP clients using the CLI. A session with the **Q** flag indicates downstream real-time analysis and that with the **u** flag indicates upstream real-time analysis:

```
(host) [mm] #show datapath session table 10.16.4.71 | include 10.16.4.80

Datapath Session Table Entries
------------------------------
Flags: F - fast age, S - src NAT, N - dest NAT
D - deny, R - redirect, Y - no syn
H - high prio, P - set prio, T - set ToS
C - client, M - mirror, V - VOIP
Q - Real-Time Quality analysis
u - Upstream Real-Time Quality analysis
I - Deep inspect, U - Locally destined
E - Media Deep Inspect, G - media signal
r - Route Nexthop, h - High Value
B - Permanent, O - Openflow
L - Log
```

```
Source IP       Destination IP  Prot SPort DPort Cntr     Prio ToS Age Destination TAge
--------------- --------------- ---- ----- ----- -------- ---- --- --- ----------- ----
10.16.4.80      10.16.4.71      17   20008 20038  0/0      6    46  0   local       24

Packets    Bytes      Flags
---------  ---------- ----------
909        115732     HPTCIQuVBO
```

RTP Analysis Limitations

- In case of split-tunnel forwarding mode, upstream UCC score is not calculated if the calling and called party are behind the same remote AP.
- UCC score, jitter, delay, and packet loss is calculated for voice RTP streams only. These metrics are not available for video streams.
- Upstream UCC score is not supported for desktop sharing , file transfer, Apple Facetime, Spectralink SVP, Vocera, and Wi-Fi Calling ALGs.

AppRF Integration with ALGs and User Role

The QOSMOS engine does not recognize many of the UCC applications. For the ones it recognizes, it does not maintain the state of the application. Due to this limitation, it cannot provide granular visibility into the UCC applications. To resolve this limitation, in ArubaOS 8.x, voice ALGs identify the application type for supported UCC applications, so that the administrator can now use AppRF rules to deny, permit, apply QoS, or rate limit UCC application traffic. The UCC application identifies all the supported applications listed below. The UCC application identifies the application type corresponding to a media session and programs the datapath with the application ID and the priority values. It is mandatory to add the ACLs to permit specific application traffics if an ACL rule is not present in the user-role to permit RTP or RTCP traffic. Following is a list of UCC applications that can be used to create application ACLs.

Table 249: *UCC Application ACLs*

UCC Application ACL
alg-facetime
alg-ftp
alg-h323
alg-jabber-audio
alg-jabber-desktop-sharing
alg-jabber-video
alg-noe
alg-rtp
alg-sccp
alg-sip
alg-sip-audio

Table 249: *UCC Application ACLs*

UCC Application ACL
alg-sip-video
alg-skype4b-app-sharing
alg-skype4b-audio
alg-skype4b-desktop-sharing
alg-skype4b-file-transfer
alg-skype4b-secure
alg-skype4b-video
alg-svp
alg-vocera
alg-wifi-calling



ArubaOS supports Qosmos SDK version 5.3.

The following pre-defined ACL are available by default. The administrator can either add the entire ACL to the appropriate user-role or selectively use the application IDs in another ACL and add that to the appropriate user-role.

```
ip access-list session voip-applications-acl
   any any app alg-skype4b-audio permit
   any any app alg-skype4b-video permit
   any any app alg-skype4b-desktop-sharing permit
   any any app alg-skype4b-app-sharing permit
   any any app alg-sip-audio permit
   any any app alg-sip-video permit
   any any app alg-sccp permit
   any any app alg-vocera permit
   any any app alg-noe permit
   any any app alg-h323 permit
   any any app alg-jabber-audio permit
   any any app alg-jabber-video permit
   any any app alg-facetime permit
   any any app alg-wifi-calling permit
   any any app alg-rtp permit
```

The ordering of the UCC application ACE is not important except for the last ACE – **any any app alg-rtp permit**. The use of this ACE is to permit RTP traffic. This is important in a deployment having media application that is not identified by the UCC application. In such a case, the UCC application falls back to the **alg-rtp** ACE as the default application ID. If permitting random RTP traffic is a requirement, this ACE should be included in the ACL. In addition, this ACE should always be the last entry in the ACL.

An example of the ACE entries of **voip-applications-acl** follows:

```
(host) [mynode] #show ip access-list voip-applications-acl

ip access-list session voip-applications-acl
voip-applications-acl
----------------------
```

```
Priority  Source  Destination  Service  Application                      Action
TimeRange
--------  ------  -----------  -------  -----------                      ------  --------
-
1         any     any                   app alg-skype4b-audio            permit
2         any     any                   app alg-skype4b-video            permit
3         any     any                   app alg-skype4b-desktop-sharing  permit
4         any     any                   app alg-skype4b-app-sharing      permit
5         any     any                   app alg-sip-audio                permit
6         any     any                   app alg-sip-video                permit
7         any     any                   app alg-sccp                     permit
8         any     any                   app alg-vocera                   permit
9         any     any                   app alg-noe                      permit
10        any     any                   app alg-h323                     permit
11        any     any                   app alg-jabber-audio             permit
12        any     any                   app alg-facetime                 permit
13        any     any                   app alg-wifi-calling             permit
14        any     any                   app alg-rtp                      permit

Log  Expired  Queue  TOS  8021P  Blacklist  Mirror  DisScan  IPv4/6  Contract
---  -------  -----  ---  -----  ---------  ------  -------  ------  --------
              Low                                              4
              Low                                              4
              Low                                              4
              Low                                              4
              Low                                              4
              Low                                              4
              Low                                              4
              Low                                              4
              Low                                              4
              Low                                              4
              Low                                              4
              Low                                              4
              Low                                              4
              Low                                              4
```

The following example adds an ACL to permit Lync or Skype for Business audio and video traffic and reference it to the user-role:

```
(host) [mynode] (config) #ip access-list session apprf-skype4b-sacl
(host) ^[mynode] (config-submode)#any any app alg-skype4b-audio permit
(host) ^[mynode] (config-submode)#any any app alg-skype4b-video permit
```

Other ACL rules like bandwidth contract, deny, 802.1p priority, and ToS can be used along with the ACL application IDs.

Microsoft Lync or Skype for Business

ArubaOS provides a seamless user experience for Microsoft® Lync or Skype for Business users using voice or video calls, app-sharing, and file-transfer in a wireless environment. Microsoft® Lync or Skype for Business is an enterprise solution for UCC. It provides support for voice, video, app-sharing, and file-transfer. The Lync or Skype for Business SDN API provides an interface to Mobility Master to access Lync or Skype for Business network diagnostic information about voice, video, app-sharing, and file-transfer without having to see into the traffic.

Lync or Skype for Business Media Classification Support in Mobility Master

By default, all the VoIP traffic undergo Media Classification on the managed device whenever RTP Traffic reaches the managed device. UCM in Mobility Master can identify and prioritize calls made using Lync or Skype for Business ALG. UCM also provides visibility for all voice calls made using the Lync or Skype for

Business ALG. UCM on Mobility Master dynamically opens firewall ports for voice and video traffic. The user does not have to explicitly define a firewall policy to permit such traffic.

The following sections describe:

- UCC Score for Lync or Skype for Business Media Classification on page 1103
- Available Call Quality Metrics on page 1103
- Lync or Skype for Business Media Classification Limitations on page 1103

UCC Score for Lync or Skype for Business Media Classification

ArubaOS supports UCC score for Lync or Skype for Business calls prioritized using media classification. As part of this feature, UCM supports the following:

- Real-time quality analysis for Lync or Skype for Business voice and video calls (voice RTP streams only)
- Real-time computation of UCC score (delay, jitter, and packet loss) for Lync or Skype for Business VoIP calls prioritized using media classification. The UCC score is computed by the AP in the downstream direction and also at the managed device in the upstream direction.
- Call Quality vs. Client Health chart in the UCC dashboard of Mobility Master.



When VoIP calls are prioritized using media classification, end-to-end call quality metrics such as mean opinion score, delay, jitter, and packet loss are not available.

UCC score computes the quality of voice calls. It takes delay, jitter, and packet loss of Real-time Transport Protocol (RTP) packets into account. UCC score is computed on a scale of 0 to 100. To compute the UCC score, you must enable RTP Analysis on Mobility Master. For more information, see Lync or Skype for Business Media Classification Support in Mobility Master on page 1102.

Available Call Quality Metrics

Following call quality metrics are available for Lync or Skype for Business calls prioritized by media classification:

Client IP, Client Mac, ALG, Duration(approximate), Orig time(approximate), Status, Reason, Call Type (voice or video), Client Health, UCC Score, UCC Band, Source port, Destination port, Originated and modified DSCP, and WMM values.

As the RTP packets are encrypted, following call quality metrics are not available for Lync or Skype for Business calls prioritized by media classification:

Client Name, Direction, Called to, MOS, MOS band, End-to-end Delay, jitter and packet loss.



File transfer and desktop sharing sessions are not prioritized by media classification. Upstream and downstream delay, jitter, and packet loss are not available for video sessions.

The **show ucc** commands displays statistics for media classification based Lync or Skype for Business ALG. For more information on the list of commands, see the *ArubaOS Command-Line Interface Reference Guide*. The UCC dashboard displays statistics for media classification based Lync or Skype for Business ALG. For more information on UCC dashboard, see Lync or Skype for Business Media Classification Support in Mobility Master on page 1102.

Lync or Skype for Business Media Classification Limitations

- The media classification logic is applicable only for UDP-based RTP traffic, which applies to real-time voice and video calls.

- Lync or Skype for Business app-sharing and file-transfer sessions are not identified and prioritized by media classification.
- When using media classification, UCC score, jitter, delay, and packet loss is calculated only for voice RTP streams. These metrics are not available for video streams.
- Media classification does not work in split-tunnel forwarding mode.
- When VoIP calls are prioritized using media classification, end-to-end call quality metrics such as Mean Opinion Score (MOS), delay, jitter, and packet loss are not available.
- Media classification does not work when the managed device is performing a network address translation for media traffic. Media classification continues to work if the media traffic is subjected to Network Address Translation (NAT) beyond the managed device.

Lync or Skype for Business SDN API Support in Mobility Master

To take advantage of UCC Lync or Skype for Business ALG, it is recommended to use the Lync or Skype for Business SDN API. ArubaOS 8.0.0.0 supports Lync or Skype for Business SDN API 2.0, 2.1.1, 2.2, and 2.4.1. Lync or Skype for Business SDN API works with Microsoft Lync or Skype for Business server to export details about voice or video calls, app-sharing, and file-transfer to Mobility Master. The communication between the Lync or Skype for Business SDN API and Mobility Master occurs over HTTP or HTTPS.

In earlier ArubaOS versions, Lync or Skype for Business SDN Manager sends the call information messages like start of call, interim update, and end of call to all the preconfigured local controllers though the clients are not present on the respective local controller. In ArubaOS 8.x, Lync or Skype for Business SDN Manager sends this information to Mobility Master only and not the managed devices. This reduces the network traffic originating from the Lync or Skype for Business SDN Manager and also relieves the managed devices of processing unwanted call information originating from the Lync or Skype for Business SDN Manager.

Lync or Skype for Business SDN API Configuration

The Lync or Skype for Business ALG should be configured from the **/mm** node hierarchy of Mobility Master. The ALG is enabled by default.

The following procedure describes how to configure the Lync or Skype for Business ALG:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles > All Profiles**.
2. Expand **UCC** and click **Skype4B ALG Configuration**.
3. In the **Skype4B ALG Configuration** section, configure the settings described in Table 250.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Table 250: *Lync or Skype for Business ALG Configuration Parameters*

Parameter	Description
Skype4B ALG Support	Enables the Microsoft® Lync or Skype for Business ALG. The ALG is enabled by default.
Skype4B SDN Over http/https	You can set the Lync or Skype for Business SDN listen protocol over HTTP or HTTPS. On configuring this, the Lync diagnostic information will be received over HTTP or HTTPS. The default value is HTTP.

Table 250: *Lync or Skype for Business ALG Configuration Parameters*

Parameter	Description
Voice Priority	Configures the DSCP value for the voice session. The default value is 46.
Video Priority	Configures the DSCP value for the video session. The default value is 34.
App-sharing Priority	Configures the DSCP value for the app-sharing session. The default value is 34.

The following CLI commands configure the Lync or Skype for Business ALG:

```
(host) [mm] (config) #ucc skype4b
(host) ^[mm] (Skype4B ALG Configuration) #enable
(host) ^[mm] (Skype4B ALG Configuration) #priority {app-sharing <app-sharing>| video
<video>|voice <voice>}
(host) ^[mm] (Skype4B ALG Configuration) #sdn {http|https}
(host) ^[mm] (Skype4B ALG Configuration) #write memory
```

The following CLI commands display the Lync or Skype for Business ALG configuration:

```
(host) [mynode] #show ucc skype4b

Sat Jun 25 03:25:43.429 2016

Skype4B ALG Configuration
-------------------------
Parameter                   Value    Set
---------                   -----    ---
Skype4B ALG Support         Enabled
Skype4B SDN Over http/https  https
voice priority              46
video priority              34
app-sharing priority        34
```

Lync or Skype for Business SDN Manager Configuration

Lync Dialog Listener must be installed and configured on the Lync front-end server. Lync or Skype for Business SDN Manager must be installed on a separate Windows 2008 or 2012 server (not on the Lync front-end server). If there are multiple front-end servers, Lync Dialog Listener should be installed on each server and configured to point at the Lync or Skype for Business SDN Manager. On Lync SDN Manager, the Mobility Master information needs to be configured.

Depending on the transport mode configured in the Lync or Skype for Business SDN Manager, the same transport mode (HTTP or HTTPS) should be configured in the **Configuration > System > Profiles > All Profiles > UCC > Skype4B ALG Configuration** page of WebUI. The following configuration is a snippet of the Lync or Skype for Business SDN Manager configuration:



Lync/Skype for Business SDN Manager is a third-party product. The following configuration is an example and provided for illustration purposes only. If you plan to use this sample in your environment, ensure that the sample meets your IT guidelines. By running this sample configuration, you acknowledge that Aruba is in no way liable for any loss, damage, or problems arising from running this sample configuration.

Read and follow the installation instruction PDF that comes with the SDN API software to install Lync or Skype for Business Dialog Listener and Lync or Skype for Business SDN Manager.

```
<Configuration Version="2.0" culture="en-US" Kind="Subscriber" Identifier="Aruba"
LastModified="2015-10-27T13:06:59.7745572Z">
<parameter key="submituri">http://10.15.16.123:32000</parameter>
<parameter key="outputschema">D</parameter>
```

```
<parameter key="clientcertificateid"></parameter>
<parameter key="domainfilters"></parameter>
<parameter key="subnetfilters"></parameter>
```

As displayed in the above configuration, the Mobility Master IP address is added to the Lync or Skype for Business SDN Manager instead of the managed device IP address and the port number has to be 32000 which is a fixed port and not a configurable parameter on Mobility Master. The general format of the submit Uniform Resource Identifier (URI) is as follows:

`http[s]://<Mobility Master-IP or fqdn>:32000`

IP Session ACL and User Role Configuration

The following procedure configures a user-role for Lync or Skype for Business clients. In addition, the procedure provides steps to add an ACL to permit TCP traffic for app-sharing and file-transfer sessions.

The following procedure describes how to configure the IP session ACL to permit TCP traffic for app-sharing and file-transfer sessions:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Roles & Policies**.
2. In the **Roles** tab, select an existing role.
3. In the **Roles > <custom-role>** section, click **Show Advanced View**.
4. In the **Policies** tab, click **+**. The **Add Policy** pop-up window opens.
5. In the **Add Policy** window, select the **Add existing session policy** option.
6. In the **Policy Name** drop-down list, select the **skype4b-acl** policy.
7. In the **Policy type** drop-down list, select **Session**.
8. Click **Submit**.
9. Repeat steps 4 and 5.
10. In the **Policy Name** drop-down list, select the **voip-application-acl** policy.
11. In the **Policy type** drop-down list, select **Session**.
12. Click **Submit**.
13. Click **Pending Changes**.
14. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the IP session ACL to permit TCP traffic for app-sharing and file-transfer sessions:

```
(host) [md] (config) #user-role  S4B-role
(host) ^[md] (config-submode) #session-acl skype4b-acl
(host) ^[md] (config-submode) #session-acl voip-applications-acl
(host) ^[md] (config-submode) #write memory
```

Lync or Skype for Business Troubleshooting

The following section describes the step-by-step procedure to troubleshoot Lync or Skype for Business ALG:

1. Ensure that the global prerequisites to enable UCC in ArubaOS 8.x is configured. For more information, see Lync or Skype for Business Troubleshooting on page 1106.
2. Connect clients to the SSID; launch the Lync 2010 or 2013 or Skype for Business application; and make audio and video calls between them.
3. Make a few calls between clients. Execute the **show ucc client-info** and **show ucc call-info cdrs** commands and also access the UCC dashboard on the WebUI to view Lync or Skype for Business call statistics and prioritization.

   ```
   (host) [mm] #show ucc client-info
   Client Status:
   ```

```
--------------
Client IP   Client MAC         Client Name  ALG      Server(IP)  Registration
State
---------   ----------         -----------  ---      ----------  ----------------
--
10.16.4.76  00:24:d7:40:c0:a0  Derek        Skype4B              REGISTERED

10.16.4.71  00:21:6b:9d:f2:74  Allen        Skype4B              REGISTERED


Call Status  AP Name  Flags  Device Type  Home_Agent  Foreign_Agent
-----------  -------  -----  -----------  ----------  -------------
In-Call      2_205           Win 7        10.16.4.9   NA
In-Call      2_205           Win 7        10.16.4.9   NA

Total Client Entries:2
Flags: V - Visitor, A - Away, W - Wired, R - Remote, E - External

(host) [mm] #show ucc call-info cdrs

Help:   [C] - Metric calculated at the Controller
        [A] - Metric calculated at the AP
CDR:
----
CDR ID  UCC Call ID  Client IP   Client MAC         Client Name  ALG
------  -----------  ---------   ----------         -----------  ---
4       2            10.16.4.71  00:21:6b:9d:f2:74  Derek        Skype4B
3       2            10.16.4.76  00:24:d7:40:c0:a0  Allen        Skype4B
1       NA           10.15.12.86 fc:c2:de:6c:01:9c  NA           WiFi-Calling

Dir  Called to  Dur(sec)  Orig Time        Status  Reason      Call Type
---  ---------  --------  ---------        ------  ------      ---------
OG   Scott      6         Nov 27 08:44:45  ACTIVE  NA          Voice
IC   Scott      6         Nov 27 08:44:45  ACTIVE  NA          Voice
NA   NA         88        Jun 4 06:41:40   ACTIVE  NA          Voice

Client Health  UCC Score[C]  UCC Score[A]  MOS         Server(IP)
-------------  ------------  ------------  ---         -----------
80             70.80/Good    38.50/Fair    4.10/Good
85             77.88/Good    41.53/Fair    4.32/Good
93             NA            NA            NA          T-Mobile

Total Entries: 2
```



The UCC Call ID, Client Name, Dir, Called to, UCC Score[C], UCC Score[A], and MOS values are not available for Wi-Fi Calling calls. The Server (IP) value displays the name of the service provider for WiFi-calling calls.

4. Execute the **show datapath session table** command on the managed device to verify if the calls are prioritized. A client with the **Q** flag indicates real-time analysis and a client with **u** flag indicates RTP analysis of upstream VoIP calls.

```
(host-mn) #show datapath session table 10.16.4.67
Datapath Session Table Entries
------------------------------

Flags: F - fast age, S - src NAT, N - dest NAT
D - deny, R - redirect, Y - no syn
H - high prio, P - set prio, T - set ToS
```

```
C - client, M - mirror, V - VOIP
Q - Real-Time Quality analysis
u - Upstream Real-Time Quality analysis
I - Deep inspect, U - Locally destined
E - Media Deep Inspect, G - media signal
r - Route Nexthop, h - High Value
B - Permanent, O - Openflow
L - Log

Source IP       Destination IP  Prot SPort DPort Cntr     Prio ToS Age
--------------- --------------- ---- ----- ----- -------- ---- --- ---
10.16.4.72      10.16.4.67      17   20002 20008  0/0     5    40  0
10.16.4.72      10.16.4.67      17   20003 20008  0/0     0    0   1
10.16.4.67      10.16.4.72      17   20012 20039  0/0     6    46  0

Destination TAge Packets    Bytes     Flags
----------- ---- ---------  --------- -----------
tunnel 22   24   398        151870    FHPTVQuI
tunnel 10   23   2          252       FCIE
local       f2   0          0         FYHPTMCVBO
```

5. Execute the **show ucc client-info** command to verify if the Lync or Skype for Business clients are in **In-Call** state.

```
(host) [mm] #show ucc client-info
Client Status:
--------------
Client IP   Client MAC         Client Name  ALG      Server(IP)  Registration
State
---------   ----------         -----------  ---      ----------  ----------------
--
10.16.4.76  00:24:d7:40:c0:a0  Derek        Skype4B              REGISTERED

10.16.4.71  00:21:6b:9d:f2:74  Allen        Skype4B              REGISTERED


Call Status  AP Name  Flags  Device Type  Home_Agent  Foreign_Agent
-----------  -------  -----  -----------  ----------  -------------
In-Call      2_205           Win 7        10.16.4.9   NA
In-Call      2_205           Win 7        10.16.4.9   NA

Total Client Entries:2
Flags: V - Visitor, A - Away, W - Wired, R - Remote, E - External
```

6. If Mobility Master uses Lync or Skype for Business SDN API, you can view the client name, called party, and end to end call quality for every Lync or Skype for Business calls on Mobility Master.
7. Execute the **show ucc trace-buffer skype4b** command to verify if Mobility Master is receiving and processing call information from the Lync or Skype for Business SDN manager.

```
(host) [mm] #show ucc trace-buffer skype4B
Skype4b Voice Client(s) Message Trace
-------------------------------------
Client IP     Client MAC         Client Name  Direction  Event Time       BSSID
---------     ----------         -----------  ---------  ----------       -----
192.0.2.22    00:23:33:41:c8:b8  Alex         OG         Jan  3 11:24:34
9c:1c:12:8a:b5:50
192.0.2.26    24:77:03:9a:6c:dc  John         OG         Jan  3 11:24:34
9c:1c:12:8a:b5:50
192.0.2.29    00:22:90:ea:9e:f1  Steve        OG         Jan  3 11:24:08
9c:1c:12:8a:b5:50

Called To  Media Type   AP Name  Src Port     Dest Port    Call Status
```

```
---------  ----------   -------  --------     ---------    -----------
Joe        Voice/Video  AP-225   50030/58008  50032/58006  Start of call
Mike       Voice/Video  AP-225   50032/58006  50030/58008  InCallQuality Update
Ken        Voice        AP-225   50026        50038        Call Quality Update

Num of Rows:3
```

8. Execute the **show datapath session table** command on the managed device and look for the **O** flag indicating if the RTP or RTCP-related flows is installed on the managed device using OpenFlow protocol.

```
(host-mn) #show datapath session table 10.16.4.67
Datapath Session Table Entries
------------------------------

Flags: F - fast age, S - src NAT, N - dest NAT
D - deny, R - redirect, Y - no syn
H - high prio, P - set prio, T - set ToS
C - client, M - mirror, V - VOIP
Q - Real-Time Quality analysis
u - Upstream Real-Time Quality analysis
I - Deep inspect, U - Locally destined
E - Media Deep Inspect, G - media signal
r - Route Nexthop, h - High Value
B - Permanent, O - Openflow
L - Log

Source IP       Destination IP  Prot SPort DPort Cntr     Prio ToS Age
--------------- --------------- ---- ----- ----- -------- ---- --- ---
10.15.17.207    10.15.17.202    17   20004 20005  0/0     6    46  0
10.15.17.202    10.15.17.207    17   20005 20004  0/0     6    46  0

Destination TAge Packets    Bytes     Flags
----------- ---- ---------  --------- -----------
local       1b   325        46131     FHPTCIVBO
local       3    347        41474     FHPTCIVBO
```

9. Execute the **show ucc call-info cdrs** to display the Call Detail Record (CDR) information or access the **Dashboard > UCC** page on the WebUI to verify if the calls are identified and prioritized.

```
(host) [mm] #show ucc call-info cdrs
Help:   [C] - Metric calculated at the Controller
        [A] - Metric calculated at the AP
CDR:
----
CDR ID  UCC Call ID  Client IP   Client MAC         Client Name  ALG
------  -----------  ---------   ----------         -----------  ---
4       2            10.16.4.71  00:21:6b:9d:f2:74  Derek        Skype4B
3       2            10.16.4.76  00:24:d7:40:c0:a0  Allen        Skype4B
1       NA           10.15.12.86 fc:c2:de:6c:01:9c  NA           WiFi-Calling

Dir  Called to  Dur(sec)  Orig Time        Status  Reason      Call Type
---  ---------  --------  ---------        ------  ------      ---------
OG   Scott      6         Nov 27 08:44:45  ACTIVE  NA          Voice
IC   Scott      6         Nov 27 08:44:45  ACTIVE  NA          Voice
NA   NA         88        Jun 4 06:41:40   ACTIVE  NA          Voice

Client Health  UCC Score[C]  UCC Score[A]  MOS         Server(IP)
-------------  ------------  ------------  ---         ----------
80             70.80/Good    38.50/Fair    4.10/Good
85             77.88/Good    41.53/Fair    4.32/Good
93             NA            NA            NA          T-Mobile
```

```
Total Entries: 2
```



The UCC Call ID, Client Name, Dir, Called to, UCC Score[C], UCC Score[A], and MOS values are not available for Wi-Fi Calling calls. The Server (IP) value displays the name of the service provider for WiFi-calling calls.

10. Execute the **show ucc statistics counter call client** and **show ucc statistics counter call global** commands to view the different call metrics.

```
(host) [mm] #show ucc statistics counter call client
Per Client Call Counters:
-------------------------
Client IP   Client MAC         Call Originated  Call Terminated  Active  Success
---------   ----------         ---------------  ---------------  ------  -------
10.16.4.76  00:24:d7:40:c0:a0  0                1                1       0
10.16.4.71  00:21:6b:9d:f2:74  0                1                1       0
10.16.4.79  00:24:d7:40:ff:a0  0                0                0       0

Failed  Blocked  Aborted  Forwarded  WMM AC-VI  WMM AC-VO  WMM-BK  WMM-BE
------  -------  -------  ---------  ---------  ---------  ------  ------
0       0        0        0          0          0          0       1
0       0        0        0          0          0          0       1
0       0        0        0          0          0          0       0

WMM (VI, VO, BK, BE):total calls with received priority

(host) [mynode] #show ucc statistics counter call global

System-wide Call Counters:
--------------------------
Call Originated  Call Terminated  Active  Success  Failed  Blocked  Aborted
---------------  ---------------  ------  -------  ------  -------  -------
0                2                2       0        0       0        0

Forwarded  WMM AC-VI  WMM AC-VO  WMM-BK  WMM-BE
---------  ---------  ---------  ------  ------
0          0          0          0       2

Device Type Allocations:
------------------------
Device Type  WMM AC-VI  WMM AC-VO  WMM-BK  WMM-BE
-----------  ---------  ---------  ------  ------
Win 7        0          0          0       2

WMM (VI, VO, BK, BE):total calls with received priority)
```

11. If the clients are not seen after executing the **show ucc client-info** command, verify the output by executing the **show gsm debug channel ucc_client**, **show gsm debug channel ucc_session**, and **show gsm debug channel ip_user**.

```
(host) [mm] #show gsm debug channel ucc_client
ucc_client Channel Table
------------------------
state  uc_client_mac      uc_client_ip  uc_contact_name  uc_server_name  uc_
client_flags
-----  -------------      ------------  ---------------  --------------  --------
-------
ACTV   80:86:f2:40:b3:d4  10.15.88.247  1008             10.15.16.30     1
```

```
ACTV   80:86:f2:40:14:9c  10.15.88.245  1007            10.15.16.30     1


uc_reg_state  uc_alg  uc_entry_type  uc_role  uc_active_call  uc_replicatorip
------------  ------  -------------  -------  --------------  ---------------
4             14      1              0        0               10.15.88.100
4             14      1              0        0               10.15.88.100

Total Num of Objects            :2
Total Num of Active Objects     :2
Total Num of Replicated Objects :0

(host) [mm] #show gsm debug channel ucc_session

ucc_session Channel Table
-------------------------
state  uc_client_mac      uc_client_ip  uc_active_call
-----  -------------      ------------  --------------
ACTV   80:86:f2:40:b3:d4  10.15.88.247  0
ACTV   80:86:f2:40:14:9c  10.15.88.245  0

Total Num of Objects            :2
Total Num of Active Objects     :2
Total Num of Replicated Objects :0

(host) [mm] #show gsm debug channel ip_user

ip_user Channel Table
---------------------
state  v_repkey  user_ip_address  user_uuid                ip_user_flags  ip_
user_timestamp
-----  --------  ---------------  ---------                -------------  ------
-----------
REPL   3         10.15.88.245     001a1e01b2280000002f0064  0
181193397240
REPL   3         10.15.88.247     001a1e01b2280000002f0065  0
181193397370

Total Num of Objects            :2
Total Num of Active Objects     :0
Total Num of Replicated Objects :2

Total number of hosts: 3
```

12. Execute the **show ucc rtpa-report** command to view the Real-Time analysis report.

```
(host) [mm] #show ucc rtpa-report
Help:   [C] - Metric calculated at the Controller
        [A] - Metric calculated at the AP
        [E] - Metric calculated End-to-End

Real-Time Analysis Call Quality Report
--------------------------------------
Client(IP)  Client(MAC)        Client(Name)  ALG      Jitter(usec)[C]
----------  -----------        ------------  ---      ---------------
10.16.4.76  00:24:d7:40:c0:a0  Derek         Skype4B  308.200
10.16.4.71  00:21:6b:9d:f2:74  Allen         Skype4B  1119.080

Pkt-loss(%)[C]  Delay(usec)[C]  UCC Score[C]  Jitter(usec)[A]  Pkt-loss(%)[A]
--------------  --------------  ------------  ---------------  --------------
0.000           118.000         92.346        36.840           0.000
0.000           35.400          76.210        101.679          0.000
```

```
Delay(usec)[A]  UCC Score[A]  Forward mode
--------------  ------------  ------------
344.610         40.116        tunnel
581.034         48.956        tunnel

Num Records:2
```

13. Execute the **show openflow-controller hosts** command to verify if the users are learned as OpenFlow hosts. If a host entry is not present for a user then flow will not be installed and the call will not be prioritized.

```
(host) [mm] #show openflow-controller hosts
Hosts
-----
IP                  MAC               Wireless  Dpid
--                  ---               --------  ----
10.15.88.245        80:86:f2:40:14:9c  True      00:00:00:1a:1e:01:b2:28
10.15.88.235        ac:bc:32:78:33:a1  True      00:00:00:1a:1e:01:b2:28
10.15.19.39         00:0c:29:e4:88:93  false     00:00:00:0c:29:e8:b8:b9

Port No  Port MAC
-------  --------
21       ac:a3:1e:ca:7d:c0
19       d8:c7:c8:c9:23:8b
1        00:0c:29:e8:b8:ba

Total number of hosts: 3
```

14. Execute the **show openflow-controller flow-table** command on Mobility Master to verify if the flows are installed accurately.

```
(host) [mm] #show openflow-controller flow-table
Flow-table
----------
Dpid                     In Port  Src Mac  Dst Mac  Ether   Src IP
----                     -------  -------  -------  -----   ------
00:00:00:0c:29:a1:de:01  *        *        *        0x800   222.173.190.239
00:00:00:0b:86:9a:16:77  *        *        *        0x800   222.173.190.239

Dst IP           Proto  Src Port  Dst Port  App Name        Actions
------           -----  --------  --------  --------        -------
186.173.202.254  17     60000     60000     ucm             output=controller
186.173.202.254  17     60000     60000     ucm             output=controller

Total number of flows: 2
```

15. Execute the **show openflow flow-table** command on the managed device to check if the OpenFlows are getting programmed in the managed device.

```
(host-mn) #show openflow flow-table
Openflow Flow Table
-------------------
In Port  Src Mac  Dst Mac  Ether   Src IP           Dst IP           Proto
-------  -------  -------  -----   ------           ------           -----
*        *        *        0x800   1.1.1.1          2.2.2.2          97
*        *        *        0x800   192.168.201.251  192.168.201.250  6
*        *        *        0x800   222.173.190.239  186.173.202.254  17
*        *        *        0x800   192.168.201.250  192.168.201.251  6


Src Port  Dst Port  Packets  Bytes    Actions
--------  --------  -------  -----    -------
```

```
*          *          0         0         (Output:controller)
42017      42008      0         0         ,(Set IP ToS:34),(Set Vlan pcp:5),(Set
AppID:2565)(Output:normal),(Write Flag:VH)
60000      60000      0         0         (Output:controller)
42008      42017      0         0         ,(Set IP ToS:34),(Set Vlan pcp:5),(Set
AppID:2565)(Output:normal),(Write Flag:VH)

Total number of flows: 4
```

Cisco Jabber

Cisco Jabber is an enterprise collaboration application that supports the following protocols:

- Voice call based on SIP signaling and media on RTP protocol
- Video call based on SIP signaling and media on RTP protocol
- Desktop-sharing based on SIP signaling and Binary Floor Control Protocol (BFCP) and media on RTP protocol
- File-transfer based on TCP protocol

Cisco Jabber is an all-in-one application and significant number of customers deploy this application in open SIP mode without encryption. As Cisco Jabber deployment continues to gain a larger foothold in the collaboration space, it is important to ensure QoS for its delay-sensitive applications such that there is no perceptible difference in the user experience between wireless and wired networks.

Cisco Jabber Support in ArubaOS

ArubaOS 8.x provides QoS and visibility for voice, video calls, and desktop-sharing sessions made using an unencrypted version of the Cisco Jabber client. UCM can uniquely identify and prioritize Cisco jabber voice, video calls, and desktop-sharing sessions.

Open SIP ALG Enhancements

Cisco Jabber is an all-in-one application, enabling a user to perform functions in addition to audio and video calls. The existing SIP ALG is enhanced to support Cisco Jabber.

Parser Logic Enhancement

The current SIP ALG parser is capable of handling audio calls. The same is extended to handle video calls, app-sharing, hold or resume calls, conference calls, and call transfer.

Two additional ports, namely TCP 5222 and TCP 8443 are added to the default **jabber-acl** IP access list. These ports are required for the clients to register to the server.

```
(host) [md] (config) #ip access-list session jabber-acl
(host) [md] (config-submode) #any any tcp 5222  permit
(host) [md] (config-submode) #any any tcp 8443  permit
```

The **jabber-acl** IP access list is included in the voice user-role by default. If the administrator chooses to use any other custom user-role, the ACL should be added manually to the custom user-role.

Identification of Cisco Jabber Clients

A new configuration setting **Jabber Server IP** is introduced where an administrator can configure the Cisco Jabber server (Cisco Unified Communication Manager and Cisco Unified Presence Manager) IP address for client identification. You can configure up to 16 such IP addresses.

Cisco Jabber Configuration

You should enable DPI on the managed device for Cisco Jabber to work. For more information on enabling DPI, see Cisco Jabber Configuration on page 1113.

The Cisco Jabber ALG should be configured from the **/mm** node hierarchy of Mobility Master. The ALG is enabled by default.

The following procedure describes how to configure the Cisco Jabber ALG:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles > All Profiles**.
2. Expand **UCC** and click **Jabber ALG Configuration**.
3. In the **Jabber ALG Configuration** section, configure the settings described in Table 251.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Table 251: *Cisco Jabber ALG Configuration Parameters*

Parameter	Description
Jabber ALG Support	Enables the Cisco Jabber ALG. The ALG is enabled by default.
Jabber Server IP	Configures the Cisco Jabber server (Cisco Unified Communication Manager and Cisco Unified Presence Server) IP address to uniquely identify Cisco Jabber clients. **NOTE:** This is a mandatory configuration setting.
Voice Priority	Configures the DSCP value for the voice session. The default value is 46.
Video Priority	Configures the DSCP value for the video session. The default value is 34.
App-sharing Priority	Configures the DSCP value for the app-sharing session. The default value is 34.

The following CLI commands configure Cisco Jabber ALG:

```
(host) [mm] (config) #ucc jabber
(host) ^[mm] (Jabber ALG Configuration) #enable
(host) ^[mm] (Jabber ALG Configuration) #priority {app-sharing <app-sharing>| video
<video>|voice <voice>}
(host) ^[mm] (Jabber ALG Configuration) #server-ip <server-ip>
(host) ^[mm] (Jabber ALG Configuration) #write memory
```

The following commands display the Cisco Jabber ALG configuration:

```
(host) [mynode] #show ucc jabber

Jabber ALG Configuration
------------------------
Parameter             Value        Set
---------             -----        ---
Jabber ALG Support    Enabled
Jabber server ip      10.15.16.30
Jabber server ip      10.15.16.31
voice priority        46
video priority        34
app-sharing priority  34
```

Cisco Jabber Troubleshooting

The following procedure describes how to troubleshoot the Cisco Jabber ALG:

1. Ensure that the global prerequisites to enable UCC in ArubaOS 8.x is configured. For more information, see Cisco Jabber Troubleshooting on page 1114.
2. Ensure that the Cisco Jabber server (Cisco Unified Communication Manager and Cisco Unified Presence Manager) IP addresses are configured under the Cisco Jabber ALG configuration.
3. Connect clients to the SSID.
4. Launch the Cisco Jabber application in the client and log in with the credentials to register with the Cisco Unified Communication Manager and Cisco Unified Presence Manager servers.
5. Execute the **show ucc client-info** command to verify if the ALG type is **Jabber** and the registration status is **REGISTERED**.

```
(host) [mm] #show ucc client-info
Thu Dec 03 08:48:09.077 2015

Client Status:
--------------
Client IP     Client MAC         Client Name  ALG     Server(IP)   Registration
State
---------     ----------         -----------  ---     ----------   --------------
----
10.15.88.235  ac:bc:32:78:33:a1  1019         Jabber  10.15.16.30  REGISTERED

10.15.88.247  80:86:f2:40:b3:d4  1008         Jabber  10.15.16.30  REGISTERED


Call Status  AP Name    Flags  Device Type  Home_Agent    Foreign_Agent
-----------  -------    -----  -----------  ----------    -------------
Idle         AP-105            OS X         10.15.88.100  NA
In-Call      AP-115            Win 7        10.15.88.100  NA

Total Client Entries:2
Flags: V - Visitor, A - Away, W - Wired, R - Remote, E - External
```

6. Start audio, video calls, and app-sharing sessions between Cisco Jabber clients.
7. Execute the **show ucc client-info** and **show ucc call-info cdrs** commands or access the **Dashboard > UCC** page on the WebUI to view Cisco Jabber call statistics and prioritization.

```
(host) [mm] #show ucc client-info
Thu Dec 03 08:48:09.077 2015

Client Status:
--------------
Client IP     Client MAC         Client Name  ALG     Server(IP)   Registration
State
---------     ----------         -----------  ---     ----------   --------------
----
10.15.88.235  ac:bc:32:78:33:a1  1019         Jabber  10.15.16.30  REGISTERED

10.15.88.247  80:86:f2:40:b3:d4  1008         Jabber  10.15.16.30  REGISTERED


Call Status  AP Name    Flags  Device Type  Home_Agent    Foreign_Agent
-----------  -------    -----  -----------  ----------    -------------
Idle         AP-105            OS X         10.15.88.100  NA
In-Call      AP-115            Win 7        10.15.88.100  NA

Total Client Entries:2
Flags: V - Visitor, A - Away, W - Wired, R - Remote, E - External

(host) [mm] #show ucc call-info cdrs
```

```
Thu Dec 03 08:48:23.827 2015

Help:   [C] - Metric calculated at the Controller
        [A] - Metric calculated at the AP
CDR:
----
CDR ID  UCC Call ID  Client IP   Client MAC         Client Name  ALG
------  -----------  ---------   ----------         -----------  ---
4       2            10.16.4.71  00:21:6b:9d:f2:74  Derek        Skype4B
3       2            10.16.4.76  00:24:d7:40:c0:a0  Allen        Skype4B
1       NA           10.15.12.86 fc:c2:de:6c:01:9c  NA           WiFi-Calling

Dir  Called to  Dur(sec)  Orig Time        Status  Reason      Call Type
---  ---------  --------  ---------        ------  ------      ---------
OG   Scott      6         Nov 27 08:44:45  ACTIVE  NA          Voice
IC   Scott      6         Nov 27 08:44:45  ACTIVE  NA          Voice
NA   NA         88        Jun 4 06:41:40   ACTIVE  NA          Voice

Client Health  UCC Score[C]  UCC Score[A]  MOS         Server(IP)
-------------  ------------  ------------  ---         ----------
80             70.80/Good    38.50/Fair    4.10/Good
85             77.88/Good    41.53/Fair    4.32/Good
93             NA            NA            NA          T-Mobile

Total Entries:2
```



The UCC Call ID, Client Name, Dir, Called to, UCC Score[C], UCC Score[A], and MOS values are not available for Wi-Fi Calling calls. The Server (IP) value displays the name of the service provider for WiFi-calling calls.

8. Execute the **show ucc client-info** command. If the **ALG** column displays **SIP** instead of **Jabber**, ensure that the Cisco Unified Communication Manager and Cisco Unified Presence Server IP addresses are added as part of the Cisco Jabber configuration in Mobility Master. In addition, verify if DPI is enabled on Mobility Master.
9. Execute the **show ucc trace-buffer jabber** command to verify if call signaling events such as establishing voice, video, desktop sharing, and file transfer are recorded.

```
(host) [mm] #show ucc trace-buffer jabber
Jabber Voice Client(s) Message Trace
-----------------------------------
Client IP     Client MAC         Client Name  Direction         Event Time
BSSID              Msg
---------     ----------         -----------  ---------         ----------
-----              ---
10.15.88.234  68:17:29:9f:b6:77  3002         Server-To-Client  Jul  4 22:48:28
ac:a3:1e:27:dc:00  200_OK
10.15.88.234  68:17:29:9f:b6:77  3002         Server-To-Client  Jul  4 22:48:27
ac:a3:1e:27:dc:00  100_TRYING
10.15.88.234  68:17:29:9f:b6:77  3002         Client-To-Server  Jul  4 22:48:27
ac:a3:1e:27:dc:00  REGISTER
10.15.88.234  68:17:29:9f:b6:77  3002         Server-To-Client  Jul  4 22:46:32
ac:a3:1e:27:dc:00  200_OK
```

10. If the clients are not seen after executing the **show ucc client-info** command, verify the output by executing the **show gsm debug channel ucc_client**, **show gsm debug channel ucc_session**, **show gsm debug channel ip_user**, and **show openflow-controller hosts**.

```
(host) [mm] #show gsm debug channel ucc_client
```

```
ucc_client Channel Table
------------------------
state  uc_client_mac      uc_client_ip  uc_contact_name  uc_server_name  uc_
client_flags
-----  -------------      ------------  ---------------  --------------  --------
-------
ACTV   80:86:f2:40:b3:d4  10.15.88.247  1008             10.15.16.30     1

ACTV   80:86:f2:40:14:9c  10.15.88.245  1007             10.15.16.30     1


uc_reg_state  uc_alg  uc_entry_type  uc_role  uc_active_call  uc_replicatorip
------------  ------  -------------  -------  --------------  ---------------
4             14      1              0        0               10.15.88.100
4             14      1              0        0               10.15.88.100

Total Num of Objects            :2
Total Num of Active Objects     :2
Total Num of Replicated Objects :0

(host) [mm] #show gsm debug channel ucc_session

ucc_session Channel Table
-------------------------
state  uc_client_mac      uc_client_ip  uc_active_call
-----  -------------      ------------  --------------
ACTV   80:86:f2:40:b3:d4  10.15.88.247  0
ACTV   80:86:f2:40:14:9c  10.15.88.245  0

Total Num of Objects            :2
Total Num of Active Objects     :2
Total Num of Replicated Objects :0

(host) [mm] #show gsm debug channel ip_user

ip_user Channel Table
---------------------
state  v_repkey  user_ip_address  user_uuid                ip_user_flags  ip_
user_timestamp
-----  --------  ---------------  ---------                -------------  ------
-----------
REPL   3         10.15.88.245     001a1e01b2280000002f0064  0
181193397240
REPL   3         10.15.88.247     001a1e01b2280000002f0065  0
181193397370

Total Num of Objects            :2
Total Num of Active Objects     :0
Total Num of Replicated Objects :2

(host) [mm] #show openflow-controller hosts

Hosts
-----
IP                  MAC               Wireless
--                  ---               --------
10.15.88.245        80:86:f2:40:14:9c  True
10.15.88.235        ac:bc:32:78:33:a1  True
10.15.19.39         00:0c:29:e4:88:93  false
```

```
Dpid                      Port No  Port MAC
----                      -------  --------
00:00:00:1a:1e:01:b2:28   21       ac:a3:1e:ca:7d:c0
00:00:00:1a:1e:01:b2:28   19       d8:c7:c8:c9:23:8b
00:00:00:0c:29:e8:b8:b9   1        00:0c:29:e8:b8:ba

Total number of hosts: 3
```

11. Execute the **show datapath session table** command on the managed device to verify if the calls are prioritized. The ToS values should be set for this session, along with other flags like **V**, **H**, **P**, **T**, **O**.

```
(host-mn) #show datapath session dpi table | include V,Age
C - client, M - mirror, V - VOIP
r - Route Nexthop, h - High Value

Source IP       Destination IP  Prot SPort DPort Cntr     Prio ToS Age
Destination TAge Packets
------------    --------------  ---- ----- ----- ----     ---- --- --- -----------
- ---- -------
10.15.89.250    10.15.89.231    17   26344 26112  0/0     5    34  0   local
 31   173
10.15.89.250    10.15.89.231    17   26345 26113  0/0     5    34  0   local
 31   76
10.15.89.250    10.15.89.231    17   23843 16767  0/0     6    46  0   local
 31   2
10.15.89.231    10.15.89.250    17   16767 23843  0/0     6    46  0   local
 31   5

Bytes      SIDX     AclVer   Int-Flag PktsDpi   UplnkVlan AppID
------     -----    ------   -------- -------   --------- ------------------------
185458     39b56    1632     2101     0         none      alg-jabber-video (2570)

4420       4b952    1632     2101     0         none      alg-jabber-video (2570)

220        87e7f    1632     2101     0         none      alg-jabber-audio (2569)

624        a5a70    1632     2125     0         none      alg-jabber-audio (2569)


AceIdx Flags           User-MAC           DpiTIdx
------ ---------       ------------------ -------
0/561  FHPTCVBO        00:00:00:00:00:00 5e
0/561  FHPTMCVBO       00:00:00:00:00:00 84
0/560  FHPTMCVBO       00:00:00:00:00:00 8d
0/560  FHPTMCVBO       00:00:00:00:00:00 c5
```

Cisco Jabber Limitations

The following are the list of limitations in Cisco Jabber:

- Visibility is not available for file transfer and a pure desktop-sharing sessions. In a pure desktop-sharing session, there is no simultaneous voice or video session going on.
- In a stand-alone or master controller deployment, visibility is not available for desktop -sharing with or without simultaneous voice or video session.
- If eXtensible Messaging and Presence Protocol (XMPP) signaling is not received for any reason before the SIP signaling from the Jabber client, the client will be identified as SIP and not Jabber.

Wi-Fi Calling

Wi-Fi calling service allows cellular users to make or receive calls using a Wi-Fi network instead of using the cellular network of the carrier. Wi-Fi calling allows users to place, receive calls, and text messages even when they are beyond a cellular coverage but having a Wi-Fi network coverage. Major carriers around the world support Wi-Fi calling service.

The following sections describe:

- Wi-Fi Calling Support in ArubaOS 8.x on page 1119
- Wi-Fi Calling Operation on page 1119
- Wi-Fi Calling Configuration on page 1119
- Wi-Fi Calling Troubleshooting on page 1120
- Wi-Fi Calling Limitations on page 1122

Wi-Fi Calling Support in ArubaOS 8.x

ArubaOS 8.x provides QoS for voice calls made using Wi-Fi calling. UCM can identify and prioritize calls made using Wi-Fi calling. UCM also provides visibility for all voice calls made using Wi-Fi calling.

Wi-Fi Calling Operation

At a high level, this is how Wi-Fi calling operates:

1. Wi-Fi Calling-capable handset initiates a DNS query to locate the evolved Packet Data Gateway (ePDG) of the carrier.
2. The handset establishes a persistent IPsec tunnel with ePDG.
3. Calls, text, and traffic for other services offered by the carrier are then carried over in this IPsec tunnel.

Some carriers use a standard FQDN format for ePDG that includes their Mobile Network Code (MNC) and Mobile Country Code (MCC). For example, T-Mobile uses ss.epdg.epc.mnc260.mcc310.pub.3gppnetwork.org. Others follow a different standard format. For example, AT&T uses epdg.epc.att.net. For a list of well known carrier DNS patterns, see Table 252.

Wi-Fi Calling Configuration

The Wi-Fi Calling ALG should be configured from the **/mm** node hierarchy of Mobility Master. This ALG is enabled by default.

The following procedure describes how to configure the Wi-Fi Calling ALG:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles > All Profiles**.
2. Expand **UCC** and click **Wi-Fi Calling ALG Configuration**.
3. In the **Wi-Fi Calling ALG Configuration** section, configure the settings described in Table 252.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Table 252: *Wi-Fi Calling ALG Configuration Parameters*

Parameter	Description
Wi-Fi Calling Support	Enables the Wi-Fi Calling ALG. The ALG is enabled by default.
Voice Priority	Configures the DSCP value for the voice session. The default value is 46.
DNS Pattern	**dns-pattern**–Configure the DNS pattern for the carrier. A maximum of 10 DNS patterns can be configured. DNS patterns for known carriers are configured by default. Default built-in patterns are: ■ 3 HK - wlan.three.com.hk ■ ATT - epdg.epc.att.net ■ Rogers - epdg.epc.mnc720.mcc302.pub.3gppnetwork.org ■ SmarTone - epdg.epc.mnc006.mcc454.pub.3gppnetwork.org ■ Sprint - primgw.vowifi2.spcsdns.net ■ T-Mobile - ss.epdg.epc.mnc260.mcc310.pub.3gppnetwork.org ■ Verizon - wo.vzwwo.com If the ePDG FQDN of the carrier does not match with the default patterns, use this option to configure the DNS pattern for the carrier. **NOTE:** The DNS IP address that Mobility Master learns for Wi-Fi Calling age out automatically, if there was no DNS query or response matching that IP for more than seven days. **service-provider**–Add the service provider name for enhanced visibility.

The following CLI commands configure the Wi-Fi Calling ALG:

```
(host) [mm] (config) #ucc wificalling
(host) ^[mm] (WiFiCalling Configuration) #enable
(host) ^[mm] (WiFiCalling Configuration) #priority voice <voice>
(host) ^[mm] (WiFiCalling Configuration) #dns-pattern <dns-pattern> service-provider
<service-provider>
(host) ^[mm] (WiFiCalling Configuration) #write memory
```

The following CLI command displays the Wi-Fi Calling ALG configuration:

```
(host) [mm] #show ucc wificalling

WiFiCalling Configuration
-------------------------
Parameter            Value     Set
---------            -----     ---
WiFiCalling Support  Enabled
voice priority       46
dns pattern          att.net   ATT
```

Wi-Fi Calling Troubleshooting

The following section describes the step-by-step procedure to troubleshoot Wi-Fi Calling ALG:

1. Ensure that the global prerequisites to enable UCC in ArubaOS 8.x is configured. For more information, see Wi-Fi Calling on page 1118.
2. Connect the 'Wi-Fi Calling'-capable handset to the SSID.
3. Add the default **wificalling-acl** and **voip-applications-acl** ACLs to the user-role. By default, these ACLs are included in the voice user-role.

4. When the handset establishes a persistent IPsec tunnel with ePDG, it displays the Wi-Fi Calling icon.
5. Execute the **show ucc dns-ip-learning** command to verify if the ePDG IP address is learned.

```
(host) [mynode] #show ucc dns-ip-learning
DNS IP Learning:
----------------
IP Address     Service Provider
----------     ----------------
208.54.85.108  T-Mobile
208.54.73.77   T-Mobile
208.54.70.110  T-Mobile
208.54.77.253  T-Mobile
208.54.75.2    T-Mobile
208.54.85.64   T-Mobile
208.54.73.76   T-Mobile
208.54.83.96   T-Mobile
208.54.85.111  T-Mobile

Total Entries:9
```

6. If the ePDG IP address is not learned, identify the FQDN of ePDG and add the DNS pattern of the carrier. FQDN may not be matching with any of the default, built-in DNS patterns.
7. Place a few calls and execute the **show ucc client-info** and **show ucc call-info cdrs** commands or access the **Dashboard > UCC** page on the WebUI to view Wi-Fi call statistics and prioritization.

```
(host) [mynode] #show ucc client-info
Client Status:
--------------
Client IP     Client MAC         Client Name  ALG           Server(IP)
Registration State
---------     ----------         -----------  ---           ----------  ---------
---------
10.15.17.208  fc:c2:de:6c:01:9c  Client       WiFi-Calling  T-Mobile
REGISTERED
10.15.17.206  d8:bb:2c:51:16:b2  Client       WiFi-Calling  T-Mobile
REGISTERED

Call Status  AP Name  Flags  Device Type  Home_Agent    Foreign_Agent
-----------  -------  -----  -----------  ----------    -------------
In-Call      4-105-2         Android      10.15.16.168  NA
In-Call      2-105-1         Apple        10.15.16.168  NA

Total Client Entries:2
Flags: V - Visitor, A - Away, W - Wired, R - Remote, E - External

(host) [mynode] #show ucc call-info cdrs

Help: [C] - Metric calculated at the Controller
      [A] - Metric calculated at the AP

CDR:
----
CDR ID  UCC Call ID  Client IP   Client MAC         Client Name  ALG
------  -----------  ---------   ----------         -----------  ---
4       2            10.16.4.71  00:21:6b:9d:f2:74  Derek        Skype4B
3       2            10.16.4.76  00:24:d7:40:c0:a0  Allen        Skype4B
1       NA           10.15.12.86 fc:c2:de:6c:01:9c  NA           WiFi-Calling

Dir  Called to  Dur(sec)  Orig Time        Status  Reason      Call Type
---  ---------  --------  ---------        ------  ------      ---------
```

```
OG    Scott       6          Nov 27 08:44:45  ACTIVE   NA           Voice
IC    Scott       6          Nov 27 08:44:45  ACTIVE   NA           Voice
NA    NA          88         Jun 4 06:41:40   ACTIVE   NA           Voice

Client Health  UCC Score[C]  UCC Score[A]  MOS          Server(IP)
-------------  ------------  ------------  ---          ----------
80             70.80/Good    38.50/Fair    4.10/Good
85             77.88/Good    41.53/Fair    4.32/Good
93             NA            NA            NA           T-Mobile

Total Entries: 2
```



The UCC Call ID, Client Name, Dir, Called to, UCC Score[C], UCC Score[A], and MOS values are not available for Wi-Fi Calling calls. The Server (IP) value displays the name of the service provider for WiFi-calling calls.

8. Execute the **show datapath session table** command on the managed device to ensure that media classification flags (**I** & **E**) are set for IPsec session destined to the ePDG IP address.
9. When a Wi-Fi Calling call is identified, the **I** and **E** flags are removed from the IPsec session and appropriate ToS and 802.1p values are set for this session, along with other flags like **V**, **H**, **P**, **T**, **O**. This occurs on the managed device.
10. When the call ends, ToS and 802.1p values are removed for the IPsec session along with the **V**, **H**, **P**, **T**, **O** flags, and the **I** and **E** flags are set. For a list of flags, execute the **show datapath session table** command on the managed device.

Wi-Fi Calling Limitations

The following is a list of limitations in Wi-Fi Calling:

- Wi-Fi Calling is not supported for clients in split-tunnel and bridge-forwarding mode.
- WLAN and end to end quality metrics are not available for Wi-Fi Calling calls.
- Wi-Fi Calling calls may get dropped in the event of a cluster failover.
- Wi-Fi Calling calls do not get prioritized when Mobility Master is not reachable. This limitation does not apply for master controller deployment.
- If a Wi-Fi Calling client roams from one managed device to another, subsequent calls may not get prioritized until the client does a DNS query for carrier ePDG.
- Wi-Fi Calling is not identified and prioritized if NAT is enabled on the user VLAN. Wi-Fi Calling is not identified and prioritized if the corresponding sessions undergo NATting by the managed device.

UCC Dashboard

The UCC dashboard gives a complete view of the UCC deployment in Mobility Master. For more information on UCC dashboard, see UCC on page 451.

UCC-AirWave Integration

The UCC-AirWave integration provides a multi-managed device visibility into the UCC solution across deployments. The Mobility Master sends raw UCC data using AMON periodically. AMP receives these AMON messages and uses this data to display user-friendly aggregated and per-client UCC statistics in AirWave. This helps the administrator to assess the overall health and troubleshoot UCC deployments in a multi-managed device environment. The UCC dashboard is supported in AirWave 8.0 onwards. You can get UCC data in AirWave from Mobility Master.

The following sections describe:

- Enabling UCC Data Collection in AirWave on page 1123
- Add AirWave as a Management Server in Mobility Master on page 1123
- Enable UCC Monitoring in Mobility Master on page 1123
- Verify the Configuration on page 1124

Enabling UCC Data Collection in AirWave

The following procedure describes how to enable UCC data collection in the AirWave WebUI:

1. Navigate to the **AMP Setup > General** tab of the AirWave WebUI.
2. In the **Additional AMP Services** section, change the **Enable UCC Data Collection** option to **Yes**.

Add AirWave as a Management Server in Mobility Master

The following procedure describes how to add AMP as a management server in Mobility Master:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Airwave**.
2. Click the **Connect to airwave** toggle switch.
3. In the **Airwave IP address** text-box, enter the AirWave server IP.
4. In the **SNMP version** drop-down list, select the appropriate version.
5. In the **Community string** drop-down list, select an existing community string or enter a new community string.
6. Click **Submit**.
7. Navigate to **Configuration > System > More > General**.
8. In **MON Receivers**, click the newly added AirWave server.
9. In **Edit MON Receiver**, enter the following detail:
    a. In the **Profile list** drop-down list, select the **default-amp** profile.
10. Click **Submit**.
11. Click **Pending Changes**.
12. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command to add AMP as a management server in Mobility Master:

```
(host)[mm] (config) #mgmt-server primary-server <primary-server-ip> profile default-amp
```

Enable UCC Monitoring in Mobility Master

By default, UCC monitoring is disabled in Mobility Master. You can enable this setting using the WebUI or CLI.

The following procedure describes how to enable UCC monitoring in Mobility Master:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In **All Profiles**, expand **Controller Profile > Mgmt Config**. Select the **default-amp** profile. This example uses the *default-amp* profile.
3. In the **Mgmt Config profile**, select the **UCC Monitoring** check box.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command to enable UCC monitoring:

```
(host) [mm] (config) #mgmt-server profile default-amp
(host) ^[mm] (Mgmt Config profile "default-amp")#uccmonitoring-enable
```

Verify the Configuration

Run the following CLI command to view the management server configuration profile:

```
(host) [mm] #show mgmt-server profile default-amp

Mgmt Config profile "default-amp" (Predefined (changed))
--------------------------------------------------------
Parameter                          Value
---------                          -----
Stats                              Enabled
Tag                                Enabled
Sessions                           Enabled
Monitored Info - Add/Update        Disabled
Monitored Info - Deletion          Disabled
Monitored Info - Periodic Snapshot Disabled
Wireless IDS Event Info            Disabled
Misc                               Enabled
Location                           Enabled
UCC Monitoring                     Enabled
AirGroup Info                      Disabled
Inline DHCP stats                  Enabled
Inline AP stats                    Enabled
Inline Auth stats                  Enabled
Inline DNS stats                   Enabled
```

Run the following CLI command to view the current Mobility Master configuration with respect to the management server configuration profile:

```
(host) [mm] #show running-config | include mgmt-server
Building Configuration...
mgmt-server  primary-server 192.0.2.1 profile default-amp
mgmt-server profile "default-ale"
mgmt-server profile "default-amp"
mgmt-server profile "default-controller"
```

The UCC-AirWave integration is complete.

UCC Limitations

- Voice ALGs are not supported when voice clients are behind a NAT device.
- Media classification does not work when user VLAN has IP NAT configured.
- When using media classification or signaling protocols, UCC score, jitter, delay, and packet loss is calculated only for voice RTP streams. These metrics are not available for video streams.
- Media classification does not work in split-tunnel forwarding mode.
- When VoIP calls are prioritized using media classification, end-to-end call quality metrics such as Mean Opinion Score (MOS), delay, jitter, and packet loss are not available.
- UCC score is calculated for voice calls and desktop-sharing sessions only.
- For Lync or Skype for Business calls, MOS is generated only for voice streams. Lync or Skype for Business server does not generate MOS for video streams, desktop-sharing, and file-transfer sessions.

Understanding Extended Voice and Video Features

This section describes the other voice and video-related functionalities that are available on Mobility Master.

Enabling WPA Fast Handover

In the 802.1X authentication profile, the WPA fast handover feature allows certain WPA clients to use a pre-authorized PMK, significantly reducing handover interruption. Check with the manufacturer of your handset to see if this feature is supported. This feature is disabled by default. The following procedure describes how to enable WPA fast handover:



This feature supports WPA clients, while opportunistic key caching (also configured in the 802.1X Authentication profile) supports WPA2 clients.

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In **All Profiles**, expand **Wireless LAN > 802.1X Authentication**. Select the **default** profile.
   This example uses the *default* profile.
3. In **802.1X Authentication Profile**, select the **WPA-Fast-Handover** check box.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure WPA fast handover:

```
(host) [md] (config) #aaa authentication dot1x default
(host) ^[md] (802.1X Authentication Profile "default") #wpa-fast-handover
```

For deployments where there are expected to be considerable delays between the managed device and APs (for example, in a remote location where an AP is not in range of another AP) you can increase the value for the bootstrap threshold in the AP System profile to minimize the chance of the AP rebooting due to temporary loss of connectivity with the managed device.

Scanning VoIP-aware ARM

ARM scanning on an AP during a call affects the voice quality. You can pause the ARM scanning on the AP when a call is active by turning on the VoIP-aware ARM scanning support to avoid voice quality issues.

The following procedure describes how to enable VoIP-aware ARM scanning in the ARM profile:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In **All Profiles**, expand **RF Management > Adaptive Radio Management (ARM)**. Select the **default-a** profile.
   This example uses the *default-a* profile.
3. In **Adaptive Radio Management (ARM) profile: default-a**, expand **Scanning** and select the **VoIP Aware Scan** check box.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

For additional information on configuring an Adaptive Radio Management profile, see Configuring ARM Profiles on page 588.

The following CLI commands enable VoIP-aware ARM scanning in the ARM profile:

```
(host) [md] (config) #rf arm-profile default-a
(host) ^[md] (Adaptive Radio Management (ARM) profile "default-a") #voip-aware-scan
```

Working with Voice over Remote Access Point

Voice traffic support is enhanced on the split-tunnel forwarding mode over a Remote AP. The voice traffic management for remote and local users are done on Mobility Master. However, the sessions are

created differently for both users. For remote users, the sessions are created on the Remote AP and for local users, the sessions are created on Mobility Master. This enhancement provides the following support for the voice traffic in the split tunnel over remote access point:

- Voice traffic QoS is consistent for both local and remote users.
- All voice ALGs work reliably in split-tunnel forwarding mode when the PBX traffic is destined to flow through the corporate network.
- Provides voice statistics and counters for remote voice clients in the split-tunnel forwarding mode.

The **Flags** parameter in the **show ucc client-info** command is updated to indicate remote users:

```
(host) [mynode] #show ucc client-info

Client Status:
--------------
Client IP      Client MAC         Client Name  ALG      Server(IP)  Registration State
Call Status
---------      ----------         -----------  -----    ----------  ------------------  --
---------
192.0.2.22     00:23:33:41:c8:b8  Alex         SIP      192.0.2.1   REGISTERED
Idle
192.0.2.26     24:77:03:9a:6c:dc  John         Jabber   192.0.2.3   REGISTERED
Idle

AP Name  Flags  Device Type  Home Agent    Foreign Agent
-------  -----  -----------  ----------    -------------
AP-105   R      OS X         192.0.2.25    NA
AP-135          Win 7        192.0.2.25    NA

Total Client Entries:2
Flags: V - Visitor, A - Away, W - Wired, R - Remote, E - External
```

Battery Boost

Battery boost is an optional feature that can be enabled for any SSIDs that support voice traffic. This feature converts all broadcast and multicast traffic to unicast before delivery to the client. Enabling battery boost on an SSID allows you to set the DTIM interval from 10 to 100 (the previous allowed values were 1 or 2), equating to 1,000 to 10,000 milliseconds. This longer interval keeps associated wireless clients from activating their radios for multicast indication and delivery, leaving them in power-save mode longer, and thus lengthening battery life. The DTIM configuration is performed on the WLAN, so no configuration is necessary on the client.

An associated parameter available on some clients is the listening interval. This defines the interval (in number of beacons) after which the client must wake to read the TIM. The TIM indicates whether there is buffered unicast traffic for each sleeping client. With battery boost enabled, the DTIM is increased but multicast traffic is buffered and delivered as unicast. Increasing the LI can further increase battery life, but can also decrease client responsiveness.



Do not enable battery boost if your network includes Polycom SpectraLink devices that use the Push-to-Talk feature.

The following procedure describes how to enable the battery boost feature and set the DTIM interval in the SSID profile:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In **All Profiles**, expand **Wireless LAN > SSID**. Select the **default** profile.
   This example uses the *default* profile.

3. In **SSID Profile**, change the **DTIM Interval** to a longer interval time.
4. Select the **Battery Boost** check box.
5. Click **Save**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure battery boost:

```
(host) [md] (config) #wlan ssid-profile defaultwlan ssid-profile <profile>
(host) ^[md] (SSID Profile "default") #battery-boost
(host) ^[md] (SSID Profile "default") #dtim-period <dtim-period>
```

Enabling LLDP

LLDP is a Layer-2 protocol that allows network devices to advertise their identity and capabilities on a LAN. Wired interfaces on Aruba APs support LLDP by periodically transmitting LLDP PDUs comprised of selected TLV elements. For a complete list of supported, see Table 253 and .

LLDP-MED is an extension to LLDP that supports interoperability between VoIP and video streaming devices and other networking clients. LLDP-MED network policy discovery lets end-points and network devices advertise the VLAN, priority levels, and DSCP values used by a voice or video application.

The following procedure describes how to configure the LLDP profile and select the TLV to be sent by the AP:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In **All Profiles**, expand **AP > AP LLDP**. Select the **default** profile.
   This example uses the *default* profile.
3. The AP LLDP profile is divided into two tabs, **General** and **Advanced**. The **General tab displays only those configuration settings that often need to be adjusted to suit a specific network. The Advanced** tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values.Both general and advanced settings are described in Table 253.

Table 253: *LLDP Profile Configuration Parameters*

Parameter	Description
General	
PDU Transmission	Select this check box to enable LLDP PDU Transmission. PDU Transmission is enabled by default.
Reception of LLDP PDUs	Select this check box to enable LLDP PDU Reception. PDU Reception is enabled by default.
Advanced	
Transmit Interval (seconds)	The interval between LLDP TLV transmission seconds. Range: 1-3600, seconds and Default: 30 seconds.
Transmit hold multiplier	The Transmit hold multiplier is a value that is multiplied by the transmit interval to determine the number of seconds to cache learned LLDP information before that information is cleared.

Parameter	Description
	If the Transmit hold multiplier value is set at its default value of 4, and the Transmit interval is at its default value of 30 seconds, then learned LLDP information will be cached for 4x30 seconds, or 120 seconds.
Optional TLVs	Select the check boxes in this section to select the optional TLV the AP interface sends in LLDP PDUs. The AP will send all optional TLV by default. ▪ **port-description:** transmit a TLV that gives a description of the AP's wired port in an alphanumeric format. ▪ **system-description**: transmit a TLV that describes the AP's model number and software version. ▪ **system-name**: transmit a TLV that sends the AP name or wired MAC address. ▪ **capabilities**: transmit the system capabilities TLV to indicate which capabilities are supported by the AP. ▪ **management-address**: transmit a TLV that indicates the AP's management IP address, in either IPv4 or IPV6 format.
802.1 TLVs	Select the check boxes in this section to select the 802.1 TLV the AP interface sends in LLDP PDUs. The AP will send all 802.1 TLV by default: ▪ **port-vlan**: transmit the LLDP 802.1 port VLAN TLV. If the native VLAN is configured on the port, the port-vlan TLV will send that value, otherwise it will send a value of "0". ▪ **vlan-name**: transmit the LLDP 802.1 VLAN name TLV. The AP sends a value of "Unknown" for VLAN 0, or "VLAN <number>" for all non-zero VLAN numbers.
LLDP-MED TLVs	Once you have associated an LLDP-MED Network policy profile with this LLDP profile, you can click the check boxes in this section to select the LLDP-MED TLV the AP interface sends in LLDP PDUs. The AP does not send any LLDP-MED TLV by default: ▪ **capabilities**: transmit the LLDP-MED capabilities TLV. The AP will automatically send this TLV if it sends any other LLDP-MED TLV. ▪ **inventory**: transmit the LLDP-MED inventory TLV. ▪ **network-policy**: transmit the LLDP-MED network-policy TLV. **NOTE:** The TLV in this section cannot be enabled unless you have associated an LLDP-MED Network policy profile

4. To associate an LLDP-MED network policy profile with the LLDP profile and the LLDP-MED TLV to be sent by the AP interface, click the **LLDP-MED network policy** that appears under the **AP LLDP > default** profile in the profile list.
5. If the LLDP profile does not currently reference an LLDP-MED Network Policy profile, you must associate an LLDP-MED Network Policy profile with the LLDP profile before you can configure any LLDP-MED settings. In **AP LLDP-MED Network Policy Profile**, click the **+** icon to link an LLDP-MED Network Policy profile.
6. Click **Submit**.

7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure LLDP:

```
(host) [md] (config) #ap lldp profile <profile>
(host) ^[md] (AP LLDP Profile "<profile>") #clone <profile>
(host) ^[md] (AP LLDP Profile "<profile>") #dot1-tlvs [port-vlan|vlan-name]
(host) ^[md] (AP LLDP Profile "<profile>") #dot3-tlvs [link-aggregation|mac|mfs|power]
(host) ^[md] (AP LLDP Profile "<profile>") #lldp-med-network-policy-profile <profile>
(host) ^[md] (AP LLDP Profile "<profile>") #lldp-med-tlvs [capabilities|inventory|network-policy]
(host) ^[md] (AP LLDP Profile "<profile>") #no ...
(host) ^[md] (AP LLDP Profile "<profile>") #optional-tlvs [capabilities|management-address|port-description|system-description|system-name]
(host) ^[md] (AP LLDP Profile "<profile>") #receive
(host) ^[md] (AP LLDP Profile "<profile>") #transmit
(host) ^[md] (AP LLDP Profile "<profile>") #transmit-hold <transmit-hold>
(host) ^[md] (AP LLDP Profile "<profile>") #transmit-interval <transmit-interval>
```

Configuring LLDP-MED Profile

The following procedure describes how to configure the LLDP-MED profile:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In **All Profiles**, expand **AP > AP LLDP-MED Network Policy**. Select the **default** profile. This example uses the *default* profile.
3. The **LLDP-MED Network Policy** profile is divided into two tabs, **General** and **Advanced**. The **General** tab displays only those configuration settings that often need to be adjusted to suit a specific network. The **Advanced** tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values. Both general and advanced settings are described in Table 254.

Table 254: *LLDP-MED Network Policy Profile Configuration Parameters*

Parameter	Description
General	
LLDP-MED application type	Click the LLDP-MED application type drop-down list and select the application type managed by this profile. ▪ **guest-voice**: if the AP services a separate voice network for guest users and visitors. ▪ **guest-voice-signaling**: if the AP is part of a network that requires a different policy for guest voice signaling than for guest voice media. Do not use this application type if both the same network policies apply to both guest voice and guest voice signaling traffic. ▪ **softphone-voice**: if the AP supports voice services using softphone software applications on devices such as PCs or laptops. ▪ **streaming-video**: if the AP supports broadcast or multicast video or other streaming video services that require specific network policy treatment. This application type is not recommended for video applications that rely on TCP with buffering. ▪ **video-conferencing**: of the AP supports video conferencing

Parameter	Description
	equipment that provides real-time, interactive video or audio services. ■ **video-signaling**: if the AP is part of a network that requires a different policy for video signaling than for the video media. Do not use this application type if both the same network policies apply to both video and video signaling traffic. ■ **voice**: if the AP services IP telephones and other appliances that support interactive voice services. This is the default application type. ■ **voice-signaling**: Select this application type if the AP is part of a network that requires a different policy for voice signaling than for the voice media. Do not use this application type if both the same network policies apply to both voice and voice signaling traffic.
LLDP-MED application VLAN	Specify a VLAN by VLAN ID (0-4094) or VLAN name.
LLDP-MED application VLAN tagging	Select this check box if the LLDP-MED policy applies to a to a VLAN that is tagged with a VLAN ID or untagged. The default value is untagged. **NOTE:** When an LLDP-MED network policy is defined for use with an untagged VLAN, then the L2 priority field is ignored and only the DSCP value is used.
Advanced	
LLDP-MED application Layer-2 priority	Specify a 802.1p priority level for the specified application type, by entering a value from 0 to 7, where 0 is the lowest priority level and 7 is the highest priority.
LLDP-MED application Differentiated Services Code Point	Select a DSCP priority value for the specified application type by specifying a value from 0 to 63, where 0 is the lowest priority level and 63 is the highest priority.

4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the LLDP-MED profile:

```
(host) [md] (config) #ap lldp med-network-policy-profile <profile>
(host) ^[md] (AP LLDP-MED Network Policy Profile "default") #application-type {guest-
voice|guest-voice-signaling|softphone-voice|streaming-video|video-conferencing|video-
signaling|voice|voice-signaling}
(host) ^[md] (AP LLDP-MED Network Policy Profile "default") #clone <profile>
(host) ^[md] (AP LLDP-MED Network Policy Profile "default") #dscp <dscp>
(host) ^[md] (AP LLDP-MED Network Policy Profile "default") #l2-priority <l2-priority>

(host) ^[md] (AP LLDP-MED Network Policy Profile "default") #no ...
(host) ^[md] (AP LLDP-MED Network Policy Profile "default") #tagged
(host) ^[md] (AP LLDP-MED Network Policy Profile "default") #vlan <vlan>
```

The following commands create a LLDP MED network policy profile for streaming video applications and marks streaming video as high-priority traffic:

```
(host) [md] (config) #ap lldp med-network-policy-profile vid-stream
(host) ^[md] (AP LLDP-MED Network Policy Profile "vid-stream") #dscp 48
```

```
(host) ^[md] (AP LLDP-MED Network Policy Profile "vid-stream") #l2-priority 6
(host) ^[md] (AP LLDP-MED Network Policy Profile "vid-stream") #tagged
(host) ^[md] (AP LLDP-MED Network Policy Profile "vid-stream") #vlan 10
```

Next, the LLDP MED network policy profile is assigned to an LLDP profile, and the LLDP profile is associated with an AP wired-port profile:

```
(host) [md] (config) #ap lldp profile video1
(host) ^[md] (AP LLDP Profile "video1") #lldp-med-network-policy-profile vid-stream
(host) ^[md] (AP LLDP Profile "video1") #!
(host) ^[md] (config)ap wired-port-profile corp2
(host) ^[md] (AP wired port profile "corp2")lldp-profile video1
```

WebRTC Prioritization

The webRTC prioritization feature prioritizes the media traffic from webRTC sources. WebRTC is an open framework for the web that enables real time communication using a web browser. WebRTC includes the fundamental building blocks for high-quality communication on the web like network, audio, and video components that are used in voice, video, and chat applications.

WebRTC prioritization provides better end user experience, dashboard visibility of all WebRTC applications like voice, video, and application sharing, and call quality monitoring for audio calls using upstream and downstream RTP analysis. Enable DPI before enabling WebRTC prioritization. To enable DPI see Enabling DPI.

WebRTC Prioritization Configuration

You can configure webRTC by using the CLI. The following commands enable webRTC on the Mobility Master using the CLI:

```
(host) [mm] (config) #ucc webrtc
(host) [mm] (WebRTC ALG Configuration) #enable
(host) [mm] (WeRTC ALG Configuration) #priority video
(host) [mm] (WeRTC ALG Configuration) #priority voice
```

WebRTC Media Classification Support in Mobility Master

By default, all the VoIP traffic undergo media classification on the managed device whenever RTP Traffic reaches the managed device. The **UCM** process in the Mobility Master can identify and prioritize calls made using webRTC ALG. The **UCM** process also provides visibility for all voice calls made using the webRTC ALG. The **UCM** process in the Mobility Master dynamically opens firewall ports for voice and video traffic. You do not have to explicitly define a firewall policy to permit such traffic.

WebRTC Media Flow Pattern

Following is the webRTC media flow pattern:

1. For any webRTC based call, there is STUN/TURN message exchanged between the webRTC endpoints.
2. The same endpoints (IP5 tuple) is used for Datagram Transport Layer Security (DTLS) packet exchange to derive the encryption/decryption keys.
3. The same endpoints (IP5 tuple) is subsequently used for Secure RTP (SRTP) packets for exchanging media.
4. The SRTP header has stream and codec details which helps in identifying the media type such as voice and video.

UCC Score for WebRTC Media Classification

ArubaOS supports UCC score for webRTC calls prioritized using media classification. The UCM process supports:

- Real-time quality analysis for webRTC voice and video calls (voice RTP streams only)
- Real-time computation of UCC score (delay, jitter, and packet loss) for webRTC VoIP calls prioritized using media classification. The AP computes the UCC score in the downstream direction and the managed device computes the UCC score in the upstream direction.
- Call quality versus client health chart is displayed in the UCC dashboard of the Mobility Master.

When VoIP calls are prioritized using media classification, end-to-end call quality metrics such as mean opinion score, delay, jitter, and packet loss are not available. UCC score computes the quality of voice calls. It takes delay, jitter, and packet loss of SRTP packets into account. UCC score is computed on a scale of 0 to 100. To compute the UCC score, enable RTP analysis on the Mobility Master.

Call Quality Metrics

Following call quality metrics are available for webRTC calls prioritized by media classification:

- Client IP address
- Client MAC address
- ALG
- Duration (approximate)
- Origination time (approximate)
- Status
- Reason
- Call type (voice or video)
- Client health
- UCC score
- UCC band
- Source port
- Destination port
- Originated and modified DSCP
- WMM values

The following call quality metrics are not available for webRTC calls prioritized by media classification because the SRTP packets are encrypted:

- Client name
- Direction
- Called to
- Mean Opinion Score (MOS)
- MOS band
- End-to-end delay
- Jitter
- Packet loss

File transfer and desktop sharing sessions are not prioritized by media classification. Upstream and downstream delay, jitter, and packet loss are not available for video sessions.

The **show ucc** commands displays statistics for media classification based WebRTC ALG. The UCC dashboard displays statistics for media classification based webRTC ALG.

WebRTC Media Classification Limitations

WebRTC media classification limitations include:

- Media classification logic is applicable only for UDP-based RTP traffic, which applies to real-time voice and video calls.
- WebRTC app-sharing and file-transfer sessions are not identified and prioritized by media classification.
- When using media classification, UCC score, jitter, delay, and packet loss are calculated only for voice RTP streams. These metrics are not available for video streams.
- Media classification does not work in split-tunnel forwarding mode.
- When VoIP calls are prioritized using media classification, end-to-end call quality metrics such as MOS, delay, jitter, and packet loss are not available.
- Media classification does not work when the managed device performs a NAT for media traffic. Media classification continues to work if the media traffic is subjected to NAT beyond the managed device.

WebRTC Troubleshooting

WebRTC allows basic troubleshooting tasks.

1. Enable UCC.
2. Connect clients to the SSID
3. Launch the webRTC application
4. Make audio and video calls between clients.
5. Issue the **show ucc client-info** and **show ucc call-info cdrs** commands and access the UCC dashboard on the webUI to view the WebRTC call statistics and prioritization

```
[mynode] #show ucc client-info
Client Status:
--------------
Client IP    Client MAC         Client Name  ALG     Server(IP)  Registration State
 Call Status  AP Name            Flags  Device Type  Home_Agent    Foreign_Agent
---------    ----------         -----------  ---     ----------  ------------------
 -----------  -------            -----  -----------  ----------    -------------
10.15.164.5  e4:b3:18:08:36:a3  Client       WebRTC              REGISTERED
 In-Call      18:64:72:c2:07:54         Windows      10.15.164.22  NA
10.15.164.7  48:51:b7:19:40:88  Client       WebRTC              REGISTERED
 In-Call      18:64:72:c2:07:54         Win 7        10.15.164.22  NA

Total Client Entries:2
Flags: V - Visitor, A - Away, W - Wired, R - Remote, E - External

(MM_UCC) [mynode] #show ucc call-info cdrs

Help: [C] - Metric calculated at the Controller
[A] - Metric calculated at the AP
CDR:
----
CDR ID  UCC Call ID  Client IP    Client MAC         Client Name  ALG
Dir  Called to  Dur(sec)  Orig Time        Status  Reason      Call Type  Client
Health  UCC Score[C]  UCC Score[A]  MOS        Server(IP)
------  -----------  ---------    ----------         -----------  ---            -
--  ---------  --------  ---------        ------  ------      ---------  ----------
---  ------------  ------------  ---        ----------
```

```
65      NA          10.15.164.7  48:51:b7:19:40:88  Client       WebRTC          NA
  NA          63        Oct 29 22:23:58  ACTIVE  NA          Video      93
 NA            NA            NA
64      NA          10.15.164.5  e4:b3:18:08:36:a3  Client       WebRTC          NA
  NA          63        Oct 29 22:23:58  ACTIVE  NA          Video      92
 NA            NA            NA
63      NA          10.15.164.5  e4:b3:18:08:36:a3  Client       WebRTC          NA
  NA          63        Oct 29 22:23:58  ACTIVE  NA          Voice      92
 69.45/Fair    84.82/Good    NA
62      NA          10.15.164.7  48:51:b7:19:40:88  Client       WebRTC          NA
  NA          63        Oct 29 22:23:58  ACTIVE  NA          Voice      93
 71.57/Good    84.61/Good    NA
```

6. Issue the **show datapath session table** command on the managed device to verify if the calls are prioritized. A client with the Q flag indicates real-time analysis and a client with u flag indicates RTP analysis of upstream VoIP calls.

```
(LOCAL1-7010) #show datapath session | include 10.15.164.7 | include 10.15.164.5
10.15.164.7       10.15.164.5     17   65525 54646  0/0     6    46  0   local
7    268        20552      FHPTCIVBOQu     6
10.15.164.7       10.15.164.5     17   65527 54648  0/0     5    34  0   local
c8   94         53711      FHPTCVBO        6
10.15.164.5       10.15.164.7     17   54648 65527  0/0     5    34  0   local
c8   117        51520      FHPTCVBO        6
10.15.164.5       10.15.164.7     17   54646 65525  0/0     6    46  0   local
2    786        71468      FHPTCIVBOQu     6
```

7. Issue the **show ucc client-info** command to verify if the webRTC clients are in InCall state.

```
(MM_UCC) [mynode] #show ucc client-info
Client Status:
--------------
Client IP    Client MAC         Client Name  ALG     Server(IP)  Registration State
Call Status  AP Name            Flags  Device Type  Home_Agent    Foreign_Agent
---------    ----------         -----------  ---     ----------  ------------------
-----------  -------            -----  -----------  ----------    -------------
10.15.164.5  e4:b3:18:08:36:a3  Client       WebRTC              REGISTERED
In-Call      18:64:72:c2:07:54         Windows      10.15.164.22  NA
10.15.164.7  48:51:b7:19:40:88  Client       WebRTC              REGISTERED
In-Call      18:64:72:c2:07:54         Win 7        10.15.164.22  NA

Total Client Entries:2
Flags: V - Visitor, A - Away, W - Wired, R - Remote, E - External
```

8. Issue the **show datapath session table** command on the managed device and look for the O flag indicating if the RTP or RTCP-related flows is installed on the managed device using OpenFlow protocol.

```
(LOCAL1-7010) #show datapath session | include 10.15.164.7 | include 10.15.164.5
10.15.164.7       10.15.164.5     17   65525 54646  0/0     6    46  0   local
7    268        20552      FHPTCIVBOQu     6
10.15.164.7       10.15.164.5     17   65527 54648  0/0     5    34  0   local
c8   94         53711      FHPTCVBO        6
10.15.164.5       10.15.164.7     17   54648 65527  0/0     5    34  0   local
c8   117        51520      FHPTCVBO        6
10.15.164.5       10.15.164.7     17   54646 65525  0/0     6    46  0   local
2    786        71468      FHPTCIVBOQu     6
```

9. Issue the **show ucc call-info cdrs** command to display the CDR information or access the **Dashboard > UCC** page in the webUI to verify if the calls are identified and prioritized.

```
(LOCAL1-7010) #show datapath session | include 10.15.164.7 | include 10.15.164.5
10.15.164.7       10.15.164.5     17   65525 54646  0/0     6    46  0   local
7    268        20552      FHPTCIVBOQu     6
```

```
10.15.164.7       10.15.164.5     17   65527 54648  0/0     5    34  0   local
 c8   94         53711      FHPTCVBO        6
10.15.164.5       10.15.164.7     17   54648 65527  0/0     5    34  0   local
 c8   117        51520      FHPTCVBO        6
10.15.164.5       10.15.164.7     17   54646 65525  0/0     6    46  0   local
 2    786        71468      FHPTCIVBOQu     6
```

10. Issue the **show ucc statistics counter call client** and **show ucc statistics counter call global** commands to view the different call metrics.

```
(MM_UCC) [mynode] #show ucc statistics counter call client
Per Client Call Counters:
-------------------------
Client IP    Client MAC         Call Originated  Call Terminated  Active  Success
Failed  Blocked  Aborted  Forwarded  WMM AC-VI  WMM AC-VO  WMM-BK  WMM-BE
---------    ----------         ---------------  ---------------  ------  -------
------  -------  -------  ---------  ---------  ---------  ------  ------
10.15.164.5  e4:b3:18:08:36:a3  22               23               2       20
23      0        0        0          10         12         0       0
10.15.164.7  48:51:b7:19:40:88  2                0                2       0
0       0        0        0          1          1          0       0

WMM (VI, VO, BK, BE):total calls with received priority

(MM_UCC) [mynode] #show ucc statistics counter call client

Per Client Call Counters:
-------------------------
Client IP    Client MAC         Call Originated  Call Terminated  Active  Success
Failed  Blocked  Aborted  Forwarded  WMM AC-VI  WMM AC-VO  WMM-BK  WMM-BE
---------    ----------         ---------------  ---------------  ------  -------
------  -------  -------  ---------  ---------  ---------  ------  ------
10.15.164.5  e4:b3:18:08:36:a3  22               23               2       20
23      0        0        0          10         12         0       0
10.15.164.7  48:51:b7:19:40:88  2                0                2       0
0       0        0        0          1          1          0       0

WMM (VI, VO, BK, BE):total calls with received priority

(MM_UCC) [mynode] #show ucc statistics counter call global

System-wide Call Counters:
--------------------------
Call Originated  Call Terminated  Active  Success  Failed  Blocked  Aborted
Forwarded  WMM AC-VI  WMM AC-VO  WMM-BK  WMM-BE
---------------  ---------------  ------  -------  ------  -------  -------  ------
---  ---------  ---------  ------  ------
42               23               4       38       23      0        0        0
    20         22         0       0

Device Type Allocations:
------------------------
Device Type  WMM AC-VI  WMM AC-VO  WMM-BK  WMM-BE
-----------  ---------  ---------  ------  ------
Win 7        10         10         0       0
Windows      1          3          0       0

WMM (VI, VO, BK, BE):total calls with received priority
```

11. If the clients are not seen after issuing the **show ucc client-info** command, verify the output by issuing the **show gsm debug channel ucc_client**, **show gsm debug channel ucc_session**, and **show gsm debug channel ip_user** commands:

```
(MM_UCC) [mynode] #show gsm debug channel ucc_client
ucc_client Channel Table
------------------------
state  uc_client_mac      uc_client_ip  uc_contact_name  uc_server_name  uc_client_
flags  uc_reg_state  uc_alg  uc_entry_type  uc_role  uc_active_call  uc_replicatorip
-----  -------------      ------------  ---------------  --------------  -----------
----  ------------  ------  -------------  -------  --------------  ---------------
ACTV   48:51:b7:19:40:88  10.15.164.7   Client                          1
     4             32      1              0        1               10.15.164.22
ACTV   e4:b3:18:08:36:a3  10.15.164.5   Client                          1
     4             32      1              0        1               10.15.164.22

Total Num of Objects            :2
Total Num of Active Objects     :2
Total Num of Replicated Objects :0

(MM_UCC) [mynode] #show gsm debug channel ucc_session

ucc_session Channel Table
-------------------------
state  uc_client_mac      uc_client_ip  uc_active_call
-----  -------------      ------------  --------------
ACTV   48:51:b7:19:40:88  10.15.164.7   1
ACTV   e4:b3:18:08:36:a3  10.15.164.5   1

Total Num of Objects            :2
Total Num of Active Objects     :2
Total Num of Replicated Objects :0

(MM_UCC) [mynode] #show gsm debug channel ip_user

ip_user Channel Table
---------------------
state  v_repkey  user_ip_address  user_uuid                ip_user_flags  ip_user_
timestamp
-----  --------  ---------------  ---------                -------------  ---------
--------
REPL   3         10.15.164.7      000b869b53f70000003622db  0
1572431351926000
REPL   3         10.15.164.5      000b869b53f70000003622e4  0
1572430151259000

Total Num of Objects            :2
Total Num of Active Objects     :0
Total Num of Replicated Objects :2
```

12. Issue the **show ucc rtpa-report** command to view the Real-Time analysis report.

```
(MM_UCC) [mynode] #show ucc rtpa-report
Help: [C] - Metric calculated at the Controller
[A] - Metric calculated at the AP
[E] - Metric calculated End-to-End

Real-Time Analysis Call Quality Report
--------------------------------------
Client(IP)   Client(MAC)        Client(Name)  ALG     Jitter(usec)[C]  Pkt-loss(%)
[C]  Delay(usec)[C]  UCC Score[C]  Jitter(usec)[A]  Pkt-loss(%)[A]  Delay(usec)[A]
UCC Score[A]  Forward mode
----------   -----------        ------------  ---     ---------------  -------------
-  --------------  ------------  ---------------  --------------  --------------  --
----------  ------------
```

```
10.15.164.5  e4:b3:18:08:36:a3  Client        WebRTC  4104.440         0.240
   177.432         69.494         9.221            0.619           448.391
88.653        tunnel
10.15.164.7  48:51:b7:19:40:88  Client        WebRTC  2359.181         0.000
   179.290         69.700         14.451           0.314           391.472
85.634        tunnel

Num Records:2
```

13. Issue the **show openflow-controller hosts** command to verify if the users are learned as OpenFlow hosts. If a host entry is not present for a user then flow is not installed and the call is not prioritized.

```
(MM_UCC) [mynode] #show openflow-controller hosts
Hosts
-----
IP              MAC                Wireless  Dpid                      Port No  Port
MAC
--              ---                --------  ----                      -------  -----
---
10.15.164.254  00:1a:1e:11:7f:80  false     00:00:00:0b:86:9b:5f:b7  1
00:0b:86:9b:5f:b8
10.15.164.24   00:0c:29:4d:72:3a  false     00:00:00:0b:86:9b:53:f7  1
00:0b:86:9b:53:f8
10.15.164.249  20:67:7c:d7:55:90  false     00:00:00:0b:86:9b:5f:b7  1
00:0b:86:9b:5f:b8
10.15.164.1    20:67:7c:83:a7:80  false     00:00:00:0b:86:9b:53:f7  1
00:0b:86:9b:53:f8
10.15.164.55   00:0c:29:de:58:2b  false     00:00:00:0c:29:f8:68:14  2
00:0c:29:f8:68:14
10.15.164.14   00:0c:29:d3:a0:95  false     00:00:00:0c:29:f8:68:14  2
00:0c:29:f8:68:14
10.15.164.30   00:0c:29:17:04:c5  false     00:00:00:0c:29:f8:68:14  2
00:0c:29:f8:68:14
10.15.164.23   00:0c:29:3c:62:a5  false     00:00:00:0c:29:f8:68:14  2
00:0c:29:f8:68:14
10.15.164.26   ec:eb:b8:8f:7e:06  false     00:00:00:0c:29:f8:68:14  2
00:0c:29:f8:68:14
10.15.164.3    00:0c:29:6f:4b:d9  false     00:00:00:0b:86:9b:53:f7  1
00:0b:86:9b:53:f8

Hosts
-----
IP           MAC                Wireless  Dpid                      Port No  Port
MAC
--           ---                --------  ----                      -------  -------
-
10.15.164.5  e4:b3:18:08:36:a3  True      00:00:00:0b:86:9b:53:f7  30
18:64:72:c2:07:54
10.15.164.7  48:51:b7:19:40:88  True      00:00:00:0b:86:9b:53:f7  30
18:64:72:c2:07:54

Total number of hosts: 12
```

14. Issue the **show openflow-controller flow-table** command on the Mobility Master to verify if the flows are installed accurately

```
(MM_UCC) [mynode] #show openflow-controller flow-table
Flow-table
----------
Dpid                    In Port  Src Mac  Dst Mac  Ether  Src IP   Dst IP   Proto
Src Port   Dst Port   App Name  Actions
```

```
----                    -------  -------  -------  -----  ------   ------   -----
--------   --------   --------  -------
00:00:00:0b:86:9b:53:f7  *        *        *        0x800  10.15.164.7
10.15.164.5      17     65525     54646     ucm       set-dscp=46,set-vlan-
priority=6,output=normal,write-flag=VHQuI
00:00:00:0b:86:9b:53:f7  *        *        *        0x800  10.15.164.5
10.15.164.7      17     54646     65525     ucm       set-dscp=46,set-vlan-
priority=6,output=normal,write-flag=VHQuI

Flow-table
----------
Dpid                    In Port  Src Mac  Dst Mac  Ether  Src IP       Dst IP
Proto  Src Port  Dst Port  App Name  Actions
----                    -------  -------  -------  -----  ------       ------
-----  --------  --------  --------  -------
00:00:00:0b:86:9b:53:f7  *        *        *        0x800  10.15.164.5  10.15.164.7
17     54648     65527     ucm       set-dscp=34,set-vlan-priority=5,set-
appid=3621,output=normal,write-flag=VH
00:00:00:0b:86:9b:53:f7  *        *        *        0x800  10.15.164.7  10.15.164.5
17     65527     54648     ucm       set-dscp=34,set-vlan-priority=5,set-
appid=3621,output=normal,write-flag=VH
```

15. Issue the **show openflow flow-table** command on the managed device to check if the OpenFlows are added in the managed device.

```
(LOCAL1-7010) #show openflow flow-table
Openflow Flow Table
-------------------
In Port  Src Mac  Dst Mac  Ether   Src IP           Dst IP           Proto  Src Port
  Dst Port   Packets  Bytes      Actions
-------  -------  -------  -----   ------           ------           -----  --------
  --------   -------  -----      -------
*        *        *        0x800   10.15.164.5      10.15.164.7      17     54648
  65527      8587     4055725    ,(Set IP DSCP:34),(Set Vlan pcp:5),(Set AppID:3621)
(Output:normal),(Write Flag:VH)
*        *        *        0x800   222.173.190.239  186.173.202.254  17     60000
  60000      0        0          (Output:controller)
*        *        *        0x800   10.15.164.7      10.15.164.5      17     65525
  54646      13512    1053494    ,(Set IP DSCP:46),(Set Vlan pcp:6)(Output:normal),
(Write Flag:QVHIu)
*        *        *        0x800   10.15.164.5      10.15.164.7      17     54646
  65525      36690    3328022    ,(Set IP DSCP:46),(Set Vlan pcp:6)(Output:normal),
(Write Flag:QVHIu)
*        *        *        0x800   10.15.164.7      10.15.164.5      17     65527
  54648      7011     4064747    ,(Set IP DSCP:34),(Set Vlan pcp:5),(Set AppID:3621)
(Output:normal),(Write Flag:VH)
```

Chapter 44
IoT

ArubaOS supports IoT applications through BLE. ArubaOS supports multiple transport mechanisms (HTTPS POSTs, WebSockets, UDP, and so on), payload encoding (JSON, Protobuf, and so on), payload content, periodicity of information updates, and so on.

Starting from ArubaOS 8.4.0.0, two standardized endpoints, Aruba telemetry-HTTPS and Aruba telemetry-websocket are provided. The message payload for these two standardized endpoints are constructed based on a published JSON schema and Protobuf format respectively. Additionally, ArubaOS supports two new device classes from EnOcean. With these new endpoints provide the capability of choosing multiple payload content from various device classes in the existing IoT transport profile.

The APB reports data comprised of advertisement and scan response frames from scanned IoT or BLE devices to the BLE daemon process (BD) that runs on an AP. BD classifies the incoming messages based on the signature of known device types. Advertisement and scan response frames for unknown devices are stored in raw format, that is, the payload for each frame type is stored as a hexadecimal string, along with the RSSI value. BD maintains a database of IoT devices with attributes populated for known or classified devices and complete advertisement or scan frames for unclassified devices.

BD periodically forwards the data to BLE Relay (BR) process that runs on the managed devices. It forwards different subsets of data to BR at varying intervals based on the requirements of the end application/consumer. For example, for enabling beacon management, BD sends data (corresponding to the state of Aruba bBeacons) to BR every 10 minutes at the beginning. Meridian endpoints can change this period to a different value though the response. For enabling Asset Management, BD sends a smaller message containing records of the received RSSI for every asset tag that is in the vicinity of an AP.

Starting from ArubaOS 8.4.0.0, the new endpoints, Aruba telemetry-HTTPS, Aruba telemetry-websocket, authentication URL, Meridian beacon management, Meridian asset tracking, UID-namespace filter for Eddystone beacon protocol, URL filter for Eddystone beacon protocol, and ZF tag endpoints are added. These endpoints add support for more than one payload content for each supported device class.l

You can configure multiple IoT or BLE devices for these endpoints. The configured IoT or BLE devices are sent to the endpoint in the IoT transport profile. This provides flexibility to choose the payload content along with the time interval and endpoints in the IoT transport profile. ArubaOS supports the following payload content:

- Managed beacons - Contains the attributes that are recognized by Aruba-specific APB. Typically, these attributes are communicated to the Meridian Editor.
- Managed tags - Contains the asset tag MAC address, RSSI, MAC address of observer or APB and timestamp.
- ZF tags - Contains the attributes that are recognized for the ZF tags.
- EnOcean sensor
- EnOcean switch
- Other iBeacons
- All BLE data

To support

Payload Encoding

The BD packages the payload to send to a particular endpoint.

- JSON encoding for beacon management data
- JSON encoding for telemetry-HTTPS
- Protobuf encoding for asset tag data
- Protobuf encoding for telemetry-websocket

For example, beacon management data to the Meridian editor is encoded as JSON string, whereas the asset tag data is encoded in the Protobuf format. ArubaOS supports the following payload encoding:

Payload Content

The payload content is the list of the classified devices that is sent to the endpoints.

ArubaOS supports Aruba Beacon Data, Aruba asset tag data, Eddystone, EnOcean Sensor, EnOcean Switch, other iBeacons, and ZF tag data payload profiles.

Transport Interval

The transport interval determines how often IoT or BLE data is sent from BD to BR. ArubaOS supports the following transport intervals:

- Asset tag RSSI data is sent every 4 seconds from each AP to Meridian
- Aruba beacon management data is sent every 30 minutes
- Telemetry-HTTPS is sent every 5 seconds
- Telemetry-websocket is sent every 1 second

Filter Attributes

The filter attributes are used to configure a list of UUIDs as a beacon whitelist. All beacons with a UUID in the configured list are reported and others are ignored.

Cell Size

The cell size defines a scanning scope, in meters, for the APB. The cell size does not have a default value. When the cell size is set a specific value and a beacon is observed outside the cell, the beacon is ignored.

Threshold Attributes

The threshold attributes are used to track the movement inside the cell size. If a BLE device that is inside the cell size moves a distance that is less than the change threshold, then the device status is not updated to the server. However, if the BLE device moves a distance that is more than the change threshold, then the device status is reported to the server. With the right threshold value, you can benefit from avoiding unstable event reporting caused by RSSI dithering.

Out of Range Timeout

The AP maintains a IoT or BLE device table for all the scanned tags. When a tag is within the scan range of an AP, the AP observes the tag, and the tag broadcasts periodically, the AP updates information, like RSSI, of the tag. If the tag loses its power and no longer broadcasts, then the tag is aged out. When an AP does not observe a tag for a period exceeding the timeout value (default value is 120 seconds), the AP performs the following:

- When the tag is within the cell size and not heard within timeout value, then the AP reports the tag as out-range.
- When the tag is not heard within 10 minutes, then the AP deletes the tag from the IoT BLE device table.

Vendor Filter

The vendor filter is either the vendor name or the vendor ID of the IoT device. The vendor ID is a 2-byte hexadecimal value preceding with 0x in 0xABCD format. The vendor name is a string that can be either a full vendor name (example: Aruba) or a substring of the actual vendor name (example: Aru) and can be case-insensitive. Configure the vendor filter in the IoT transport profile. The vendor filter accepts up to five combinations of vendor names or vendor IDs separated by commas, for example:

- Aruba, Favendo, HanVit, SoluM, ABB
- 0xABCD, 0xBCDE, 0xCDEF, 0xDEF0, 0xEF01
- Aruba, 0xABCD, Favendo, 0xBCDE, HanVit

If more than one vendor name or vendor ID is configured, then any of the matching vendor names or vendor IDs in the vendor filter is applied. A vendor filter is reported only if the vendor data or vendor name is not empty and matches the vendor information configured in the IoT transport profile. If the vendor field is not populated for the IoT devices, the IoT devices are reported because there is not matching vendor filter in the IoT transport profile.

IoT and Third Party Server Integration

ArubaOS 8.4.0.0 enables the integration of built-in IoT BLE messages with third party servers. This integration provides a flexible interface for users to build their own endpoint and service without Meridian support. Managed devices send messages received from APs to the endpoints.

Proxy Server

The proxy server configuration in an IoT transport profile allows you to send IoT data to a proxy server that can in turn relay the IoT data to its final destination. This is useful when you cannot establish a direct link with a server that is hosted in the cloud. The proxy server configuration includes the IP address and port number of the proxy sever and the optional username and password to log in to the proxy server.

Support for Zigbee USB Dongle

APs which do not have an integrated Zigbee radio (example: 300 Series access points) support Aruba Zigbee USB dongle and provide IoT services. The Aruba Zigbee USB dongle provides the same capabilities that are provided by the integrated Zigbee radio in some APs (example: 500 Series access points). The Aruba Zigbee USB dongle uses the Nordic nRF52840 chipset and is an AP accessory that plugs into the USB port of an AP. Based on available USB ports, an AP can support multiple Aruba Zigbee USB dongles at same time.

Support for SES-Imagotag Cloud TLS Authentication

ArubaOS allows an AP with ESL USB dongle to connect to the SES cloud by using TLS authentication. This allows you to configure and update the ESL through the SES cloud. Configure either an IP address or domain name in the SES profile in the managed device.

The AP initiates a TLS authentication with the SES cloud using an Aruba certificate. After a successful authentication, the AP and SES cloud use a session key to communicate with each other. If a ESL USB dongle is unplugged from the AP or the IP address or the domain name is deleted, the session between the AP with ESL USB dongle and the ESL cloud is terminated.

To allow SES-imagotag cloud TLS authentication, configure the SES-Imagotag server name or SES-Imagotag server IP address the AP system profile. The SES-Imagotag server name accepts an FQDN while the SES-Imagotag server IP address accepts an IP address. If both are configured, the SES-Imagotag server name takes higher priority and the SES-Imagotag server IP address does not take effect. If the SES-Imagotag server name is deleted, the SES-Imagotag server IP address takes effect. To disable SES-Imagotag, delete both the SES-Imagotag server name and SES-Imagotag server IP address.

To configure a SES-Imagotag server name:

```
(host) [mynode] (AP system profile "SES-Server-Name") #sesimagotag-esl-server
<sesImagotag-esl-server>
```

To configure a SES-Imagotag server IP address:

```
(host) [mynode] (AP system profile "SES-Server-IP-Address") #sesimagotag-esl-serverip
<sesImagotag-esl-serverip>
```

To delete a SES-Imagotag server name:

```
(host) [mynode] (AP system profile "SES-Server-Name) #no sesimagotag-esl-server
```

To delete a SES-Imagotag server IP address:

```
(host) [mynode] (AP system profile "SES-Server-IP-Address) #no sesimagotag-esl-serverip
```

Support for AmberBox Sensor

ArubaOS supports AmberBox detectors and gateways that connect to a USB port in an AP. The AP relays the traffic from the detector or gateway to the destination server.

Support for MySphera Tag

MySphera is a leading provider of BLE-based asset tracking tags and services. When a MySphera BLE tag broadcasts an advertisement, an AP obtains the RSSI information, computes the location of the tag, and relays the location information to a destination server. A new device class filter, MySphera that matches the server type Telemetry-HTTPS and Telemetry-Websocket to configure the MySphera data in the IoT transport profile.

Support for ABB Sensor

ArubaOS supports the following ABB ability smart sensors and forwards the IoT data from these sensors over Telemetry-HTTPS and Telemetry-websocket server types:

- Motor sensor
- Pump sensor
- Bearing sensor
- Ambient sensor

- ECM drive sensor
- CoMo sensor

Wi-Fi RTLS and BLE Telemetry Streams

The Wi-Fi RTLS and BLE telemetry streams in ArubaOS are combined into a single telemetry stream in the IoT transport profile. This optimizes the integration of telemetry streams with third party location engines. Three new device class filters, Wi-Fi RTLS tags, Wi-Fi associated stations, and Wi-Fi unassociated stations are introduced in the IoT transport profile. If required, define a MAC address filter for filtering the RTLS tags based on the destination address. If required, enable Wi-Fi client count to send counts instead of the actual device data.



Only aggregated client count per AP for the Meridian Beacons Management server type is sent. Device data is not split based on radio type (2.4 GHz or 5 GHz).

ZF Openmatics Support for ZF BLE Tag Communication

You can manage ZF TAGs and implement BLE location service using the third-party ZF Openmatics.

Support for IoT Southbound API

ArubaOS now supports a IoT Southbound API that allows interaction with IoT devices in a generic fashion and does not require any knowledge of the device by Aruba infrastructure. The IoT southbound API runs on top of the Telemetry-WebSocket interface without any additional configuration. The IoT southbound API works asynchronously and accepts actions at any time while the WebSocket connection is up.

Aruba Sensor Value

An Aruba AP actively scans for BLE devices in the background. An AP listens to both advertisement frames and scan response frames. An AP classifies and parses the content of both types of frames and reports BLE telemetry to subscribers.

For BLE devices that an AP does not support, server type Generic-Https and Generic-Websocket can transport the broadcasting frame payload without parsing. Aruba APs support any BLE device that conforms to Aruba BLE Telemetry Advertisement Specification. Hence, even if an AP does not ’t know anything else about a BLE device, it can parse and report telemetry and sensor information to a subscriber.

Configuring Aruba Sensor in IoT Transport Profile

The following CLI commands configure an Aruba sensor in an IoT transport profile:

```
(host) (config) #iot transportProfile Sample-Sensor
(host) (IoT Data Profile "Sample-Sensor") #deviceClassFilter aruba-sensors
```

Reporting Sensor Values

Some sensor values are reported periodically to the configured server. The reporting interval is configured in the IoT transport profile. However, some sensor values are reported immediately without

waiting for the next reporting interval. For example, events based on intrusion, fire, water level, and so on are reported immediately.

Immediate reporting is applicable for any sensor and the reporting is based on the sensor values. Following are some of the sensor values that are reported immediately based on any change:

- Alarm water
- Alarm smoke
- Alarm fire
- Alarm glassbreak
- Alarm intrusion
- Input switch
- Input rocker
- Sensor contact
- Vendor data

Immediate reporting is applicable for telemetry-HTTPS and telemetry-websocket server types configured in the IoT transport profile. Aruba IoT infrastructure supports any BLE device that conforms to Aruba BLE telemetry advertisement specification. If an information element type is not supported, but they are still part of the vendor specification, it is stored in vendor-data field of the BLE device in the BLE database. If there is any change to this vendor-data, it is reported immediately.



Vendor data is supported for EnOcean and Aruba sensors.

BluConsole

ArubaOS supports iOS BluConsole mobile application that allows a user to access the serial console of an AP over BLE. The BluConsole mobile application performs a secure exchange of keys with the BLE service ABLECONSOLE that runs on an AP and encrypts traffic over the BLE connection.

Before establishing a BLE connection, the BluConsole mobile application requests user permission to pair with an AP. Bonding is not performed until the user allows permission. User is requested permission every time the BluConsole mobile application connects to an AP.

Favendo Sensors

Favendo is a leading provider of BLE-based location services. The BLE relay process in Aruba AP reports telemetry streams to Favendo servers. To allow the AP to report telemetry streams to Favendo servers, see Configuring IoT on page 1151.

HanVit Sensors

HanVit is a leading provider of BLE asset tags. When a HanVit asset tag broadcasts advertisements, an AP receives the RSSI of the asset tag and computes its location. The AP transfers the location data of a BLE asset tag along with sensor data to a management server. A management server can subscribe to the telemetry data with configuration of an IoT profile.

Configuring HanVit Sensor in IoT Transport Profile

Aruba sensors match with the following server types:

- Meridian-Asset-Tracking
- Meridian-Beacon-Management
- Telemetry-Https
- Telemetry-Websocket

The following CLI commands configure an Aruba sensor in an IoT transport profile:

```
(host) [mynode] (config) #iot transportProfile Sample-Https
(host) [mynode] (IoT Data Profile "Sample-Https") #serverType Telemetry-Https
(host) (IoT Data Profile "Sample-Https") #deviceClassFilter aruba-sensors

(host) [mynode] (config) # iot transportProfile Sample-Websocket
(host) [mynode] (IoT Data Profile "Sample-Websocket") #serverType Telemetry-Websocket
(host) [mynode] (IoT Data Profile "Sample-Websocket") #deviceClassFilter aruba-sensors

(host) [mynode] (config) # iot transportProfile Sample-Meridian-Asset-Tracking
(host) [mynode] (IoT Data Profile "Sample-Meridian-Asset-Tracking") #serverType Meridian-
Asset-Tracking
(host) [mynode] (IoT Data Profile "Sample-Meridian-Asset-Tracking") #deviceClassFilter
aruba-sensors

(host) [mynode] (config) # iot transportProfile Sample-Meridian-Beacon-Management
(host) [mynode] (IoT Data Profile "Sample-Meridian-Beacon-Management") #serverType
Meridian-Beacon-Management
(host) [mynode] (IoT Data Profile "Sample-Meridian-Beacon-Management") #deviceClassFilter
aruba-sensors
```

SoluM ESL Gateway

SoluM is a leading provider of Electronic Shelf Label (ESL) services. An ESL tag is used to display product pricing on store shelves. Typically, ESLs are attached to the front edge of store shelves. The product pricing is automatically updated whenever the price is changed from a SoluM gateway server.

A SoluM gateway server pushes configuration information, ESL image, and SoluM gateway firmware through an Aruba AP to a SoluM gateway. A SoluM gateway is a 802.15.4 Zigbee-based USB dongle. A SoluM gateway plugs into the USB port of an Aruba AP, obtains an IP address, and transfers data from a SoluM gateway server to the ESL tags that are within the range of the SoluM gateway. The SoluM gateway works as a wired client to the ESL tags and SoluM gateway. The USB port of the AP works as a wired Ethernet port and supports bridge mode. An Ethernet port on the AP allows the SoluM gateway traffic to pass through.

In the bridge mode, the USB0 ether interface of the SoluM gateway which connects to the AP obtains an IP address (either static or dynamic). The USB ether interface which is created in the AP does not use an IP address, it only receives or sends data packets to the SoluM gateway without changing the packets.

A SoluM gateway supports hot plug in and can be plugged into an Aruba AP at any time irrespective of AP status (up or down). The USB port of the AP is recognized as CDC-ether and works as a normal wired port. After the AP boots, it recognizes the SoluM gateway. The configured wired port profile for the USB port is used to send or receive data packets.

The SoluM gateway can obtain an IP address by using either DHCP and or using a static configuration. However, if the if static address does not match with the network, the traffic from the SoluM gateway is not allowed.

The following CLI commands configure the USB port of an AP as an Ethernet port:

```
(host) [mynode] (config) #ap-group solum
(host) [mynode] (AP group "solum") #enet-usb-port-profile default
```

Hanshow ESL Gateway

Hanshow is a leading provider of ESL services. ESLs are attached to the front edge of store shelves. The product pricing is automatically updated whenever the price is changed from a Hanshow gateway server.

A Hanshow gateway server pushes configuration information, ESL image, and Hanshow gateway firmware through an Aruba AP to a Hanshow gateway. A Hanshow gateway is a USB dongle (running proprietary wireless protocol) that plugs into the USB port of an Aruba AP, obtains an IP address, and transfers data from a Hanshow gateway server to the ESL tags that are within the range of the Hanshow gateway. The Hanshow gateway works as a wired client to the ESL tags and Hanshow gateway. The USB port of the AP works as a wired Ethernet port and supports bridge mode. An Ethernet port on the AP allows the Hanshow gateway traffic to pass through.

In the bridge mode, the USB0 ether interface of the Hanshow gateway which connects to the AP obtains an IP address (either static or dynamic). The USB ether interface which is created in the AP does not use an IP address, it only receives or sends data packets to the Hanshow gateway without changing the packets.

A Hanshow gateway supports hot plug in and can be plugged into an Aruba AP at any time irrespective of AP status (up or down). The USB port of the AP is recognized as CDC-ether and works as a normal wired port. After the AP boots, it recognizes the Hanshow gateway. The configured wired port profile for the USB port is used to send or receive data packets.

The Hanshow gateway can obtain an IP address by using either DHCP and or using a static configuration. However, if the if static address does not match with the network, the traffic from the Hanshow gateway is not allowed.

To configure the USB port of an AP as an Ethernet port, run the following CLI command:

```
(host) [mynode] (config) #ap-group hanshow
(host) [mynode] (AP group "hanshow") #enet-usb-port-profile default
```

Wiliot Sensors

Wiliot is a leading provider of battery-free BLE tags. An Aruba AP forwards the BLE data from Wiliot sensors immediately after parsing the advertisements and scan responses. The BLE data can be either actual advertisements or scan responses that are received by the AP. The BLE data is reported to northbound servers over websockets.

To allow the AP to report telemetry-websocket streams using the CLI:

```
(host) [mynode] (config) #iot transportProfile sample
(host) [mynode] (IoT Data Profile "sample") #servertype telemetry-websocket
(host) [mynode] (IoT Data Profile "sample") #deviceClassFilter wiliot
```

To view the Iot transport profile for Wiliot sensors using the CLI:

```
(host) [mynode] #show iot transportProfile sample

IoT Data Profile "sample"
-------------------------
Parameter                       Value
---------                       -----
Server Type                     Telemetry-Websocket
Server URL                      N/A
Access Token                    N/A
Client Id                       N/A
Username                        N/A
Password                        N/A
Reporting interval              600
Device Class Filter             wiliot
UUID Filter                     N/A
```

```
Movement Filter               0
Cell Size Filter              0
Vendor Filter                 N/A
Age Filter                    0
Authentication URL            N/A
UID Namespace Filter          N/A
URL Filter                    N/A
Access ID                     N/A
Zigbee Socket Device Filter   N/A
RSSI Reporting Format         average
choose an environment type    office
Custom Fading Factor          20
Iot Proxy Server              N/A
Iot Proxy User                N/A
AP Group                      N/A
Send device counts only       Disabled
RTLS Destination MAC Address  N/A
Data Filter                   N/A
```

ZigBee Sockets

Zigbee Socket Device (ZSD) can be configured and applied as a filter in IoT transport. With ZSD, you may specify the source endpoint, destination endpoint, destination profile ID or destination cluster ID and the packets between the ZigBee devices and server are transmitted through the Aruba Telemetry Websocket.

To configure ZSD filter using the CLI:

```
(host) [mynode] (config) #iot transportProfile sample
(host) [mynode] (IoT Data Profile "sample") #ZSDFilter <zigbee-socket-device-profile>
```

To configure a ZigBee service profile using the CLI:

```
(host) [mynode] (config) #zigbee service-profile servicetest
(host) [mynode] (Zigbee Service Profile "servicetest") #panid auto
(host) [mynode] (Zigbee Service Profile "servicetest") #permit-joining on
(host) [mynode] (Zigbee Service Profile "servicetest") #radio-instance all
(host) [mynode] (Zigbee Service Profile "servicetest") #security disable
```

To configure a ZigBee socket device profile using the CLI:

```
(host) [mynode] (config) #zigbee socket-device-profile test
(host) [mynode] (Zigbee Socket Device Profile "test") #inbound <inbound-socket-profile>
(host) [mynode] (Zigbee Socket Device Profile "test") #outbound <outbound-socket-profile>
```

To configure a ZigBee inbound socket profile using the CLI:

```
(host) [mynode] (config) #zigbee socket-inbound-profile in
```

To configure a ZigBee outbound socket profile using the CLI:

```
(host) [mynode] (config) #zigbee socket-outbound-profile out
```

Support for Input-Filter on BLE Devices

When IoT transport profiles are configured, BLE-devices are filtered based on the IoT transport profiles which may include device class, UUID, or vendor filters. Only BLE devices that should be reported are stored in the BLE-table and data loss is avoided.

The following command enables the entry filter:

```
(host) [mynode] #ap ble-init-action ap-name <ap-name> input-filter-enable
```

The following command disables the entry filter:

```
(host) [mynode] #ap ble-init-action ap-name <ap-name> input-filter-disable
```

To view the latest filtered BLE-device, issue the following command:

```
(host) [mynode] #show ap debug ble-input-filter-stats ap-name <ap-name>
```

Data Filter

The data filter in the Iot transport profile reduces the traffic in the telemetry interfaces. The data filter is a string of numbers that map to the IoT API protobuf fields. Following is the mapping of the data filters and the protobuf fields:

Value	Description
#2	reporter
2.1	name
2.3	ipv4
2.4	ipv6
2.5	hwType
2.6	swVersion
2.7	swBuild
2.8	time
#3	reported
3.2	deviceClass
3.3	model
3.4	firmware
3.5	assetId
3.6	publicKey
3.7	lastSeen
3.9	bevent
3.10	rssi
3.11	cell
3.12	beacons
3.13	txpower
3.14	sensors
3.14.1	accelerometer
3.14.2	battery
3.14.3	temperatureC
3.14.4	humidity

Value	Description
3.14.5	voltage
3.14.6	illumination
3.14.7	motion
3.14.8	current
3.14.9	CO
3.14.10	CO2
3.14.11	VOC
3.14.12	resistance
3.14.13	pressure
3.14.14	alarm
3.14.15	contact
3.14.16	occupancy
3.14.17	mechanicalHandle
3.14.18	distance
3.14.19	capacitance
3.16	stats
3.16.1	uptime
3.16.2	adv_cnt
3.16.3	seq_nr
3.17	inputs
3.18	vendorData
3.19	vendorName
3.20	sensorTimestamp
3.21	flags
3.22	localName
3.23	identity

To configure the data filter in IoT transport profile using the CLI:

```
(host) [mynode] (config) #iot transportProfile sample
(host) [mynode] (IoT Data Profile "sample") #dataFilter <dataFilter>
```

To view the data filter configured in an Iot transport profile using the CLI:

```
(host) [mynode] #show iot transportProfile sample
```

```
IoT Data Profile "sample"
-------------------------
Parameter                     Value
---------                     -----
Server Type                   Telemetry-Websocket
Server URL                    N/A
Access Token                  N/A
Client Id                     N/A
Username                      N/A
Password                      N/A
Reporting interval            600
Device Class Filter           wiliot
UUID Filter                   N/A
Movement Filter               0
Cell Size Filter              0
Vendor Filter                 N/A
Age Filter                    0
Authentication URL            N/A
UID Namespace Filter          N/A
URL Filter                    N/A
Access ID                     N/A
Zigbee Socket Device Filter   N/A
RSSI Reporting Format         average
choose an environment type    office
Custom Fading Factor          20
Iot Proxy Server              N/A
Iot Proxy User                N/A
AP Group                      N/A
Send device counts only       Disabled
RTLS Destination MAC Address  N/A
Data Filter                   2.4
```

BLE Device Management

The IoT manager runs as a process on the Mobility Master and provides BLE device management and visibility. The IoT manager allows configuration of BLE devices, provides visibility of detected BLE devices, and categorizes the BLE devices. The IoT manager communicates with the IoT radio devices onboard an AP. This allows AP BLE configuration support for non-Meridian usage. The BLE devices are managed through a BLE service profile. When a BLE service profile is applied to an AP group , all APs in that AP group inherit the BLE service profile.

The following command configures a BLE service profile:

```
(host) [mynode] (config) #ble service-profile <service-profile>
```

The following command shows the BLE service profiles:

```
(host) [mynode] #show ble service-profile
```

The following command shows a specific BLE service profile:

```
(host) [mynode] #show ble service-profile <service-profile>
```

The following command shows the iBeacon information:

```
(host) [mynode] #show iot-manager ble-service ibeacon-info profile <profile-name>
```

The following command exports the iBeacon information of all devices in an AP group:

```
(host) [mynode] #iot-manager export ibeacon-info ap-group <ap-group>
```

The following command replaces an old AP of the specified BLE MAC address with a new AP of the specified BLE MAC address:

```
(host) [mynode] #iot-manager replace-ap old-ap-mac 01:02:03:04:05:06 new-ap-mac
aa:ab:ac:ad:ae:af
```

Configuring IoT

ArubaOS allows you to create IoT transport profile, radio profile, ZigBee service profile, ZigBee socket device profile, assign a ZigBee socket device profile to an IoT transport profile, and assign an IoT transport profile to an AP group through the WebUI and CLI.

Creating an IoT Transport Profile

The following procedure describes how to create an IoT transport profile in the WebUI:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > IoT** page.
2. In the **Transport streams** tab, click **+** and configure the following parameters:

Table 255: *IoT Transport Profile Configuration Parameters*

Parameter	Description
Name	Name of the IoT transport profile.
Server URL	URL of server to send the telemetry data.
Server Type	Type of the server that receives the telemetry data. Available options are: ■ Meridian-Beacon-Management ■ Meridian-Asset-Tracking ■ ZF-Openmatics ■ Telemetry-Https ■ Telemetry-Websocket ■ Assa-Abloy Telemetry-Https is the default server type.
Device Classes	Device class tags to filter the devices that are included in the reports. For the server type Meridian-Beacon-Management, only aruba-beacons device class is supported. For the server type Meridian-Asset-Tracking, only aruba-tags device class is supported. For the server type ZF-Openmatics, only zf-tags device class is supported. For server type Telemetry-Https or Telemetry-Websocket, following device classes are supported: ■ all device classes ■ aruba-beacons ■ aruba-tags ■ zf-tags ■ enocean-sensors ■ enocean-switches ■ ibeacon ■ eddystone ■ unclassified ■ assa-abloy ■ aruba-sensors ■ ability-smart-sensor ■ wifi-tags ■ wifi-assoc-sta ■ wifi-unassoc-sta ■ mysphera ■ sbeacon ■ wiliot

Table 255: *IoT Transport Profile Configuration Parameters*

Parameter	Description
	■ ZSD ■ serial-data ■ exposure-notification All device classes is the default device class.
Reporting interval	Reporting interval of the IoT transport stream. Valid range is 1 second to 3600 seconds and default interval is 600 seconds. Selecting **Report only device count** reports only the device count at each reporting interval.
Report only device count	Reports only the device count at each reporting interval.
RSSI reporting format	Format for reporting received signal strength indicator. Available options are: ■ Average ■ Last ■ Bulk ■ Max ■ Smooth Smooth is the default RSSI reporting format.
Environment type	Environment where the device is deployed. Available options are: ■ Office ■ Warehouse ■ Auditorium ■ Shipboard ■ Outdoor ■ Custom Office is the default environment type.
Fading factor	Fading factor in custom environment. This parameter is available only when **Enviroment type** is set to **Custom**. Valid range is 10 to 40 and default fading factor is 20.
Authentication	
Authentication	Authentication type used by device. Available options are: ■ Use credentials - Use the defined credentials. ■ Use Token - Use the token credentials. Use credentials is the default authentication type.
Authentication server URL	URL of the authentication server. This parameter is available only when **Authentication** is set to **Use credentials**.
Username	Username to use when authenticating with an authentication server. This parameter is available only when **Authentication** is set to **Use credentials**.
Password	Password to use when authenticating with an authentication server. This parameter is available only when **Authentication** is set to **Use credentials**.
Client ID	Identify to the client.

Table 255: *IoT Transport Profile Configuration Parameters*

Parameter	Description
Access token	Access token to use when authenticating using a token. This parameter is available only when **Authentication** is set to **Use Token**.
Proxy server	
Type	Type of the proxy server. Available options are: ■ Server - Use server as proxy server. ■ User - Use user as proxy server. Server is the default type of proxy server.
IP address	IP address of the proxy server. This parameter is available only when **Type** is set to **Server**.
Port	Port of the proxy server. This parameter is available only when **Type** is set to **Server**.
Username	Username to use when authenticating with a proxy server. This parameter is available only when **Type** is set to **User**.
Password	Password to use when authenticating with a proxy server. This parameter is available only when **Type** is set to **User**.
Device filters	
Report devices that are within n meters of the beacon	Reports devices that are within n meters of the beacon where n is the number of meters within the range 2 to 100.
Report devices that have moved more than n meters since last reported	Reports devices that have moved more than n meters since they were last reported where n is the number of meters within the range 2 to 30.
Report devices that have had activity in the last n minutes	Reports devices that have had activity in the last n minutes where n is the number of seconds within the range 30 to 3600 or the number of minutes within the range 1 to 60.
Report devices using following filters	Reports devices using one of the defined filters. This parameter is available only when **Device classes** is set to **iBeacon** or **Eddystone**. Available options are: UUID - This parameter is available only when **Device classes** is set to iBeacon. UIDs - This parameter is available only when **Device classes** is set to **Eddystone**. URLs - This parameter is available only when **Device classes** is set to **Eddystone**. Use **Add** to add more UUID, UIDs or URLs.
AP groups	
Available AP groups	Displays all defined AP groups. Select an AP group and use **>** to add the selected AP group to **Selected AP Groups**. Use **>>** to add all AP groups to **Selected AP Groups**.

Table 255: *IoT Transport Profile Configuration Parameters*

Parameter	Description
Selected AP Groups	Lists the AP groups that are added to the IoT transport profile. Select an AP group and use **<** to remove the selected AP group from **Selected AP Groups**. Use **<<** to remove all AP groups from **Selected AP Groups**.
Zigbee socket device profiles	
Zigbee socket device profiles	Displays all ZigBee socket device profiles. Select the required ZigBee socket device profiles. See Creating a ZigBee Socket Device Profile.

3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following command configures an IoT transport profile in the CLI:
```
(host) [mynode] (config) #iot transportProfile <iot-profile-name>
```
The following command displays the list of IoT transport profiles in the CLI:
```
(host) [mynode] #show iot transportProfile
```
The following command displays the details of an IoT transport profile in the CLI:
```
(host) [mynode] #show iot transportProfile <iot-profile-name>
```
The following command applies an IoT transport profile in the CLI:
```
(host) [mynode] (config) #iot useTransportProfile <iot-profile-name>
```

Example

To create an IoT transport profile for managing Meridian beacons, run the following CLI command:
```
(host) [mynode] (config) #iot transportProfile Sample-Beacons-Management
```
To create an IoT transport profile for tracking Meridian assets, run the following CLI command:
```
(host) [mynode] (config) #iot transportProfile Sample-Meridian-Asset-Tracking
```
To create an IoT transport profile for Aruba telemetry HTTPS, run the following CLI command:
```
(host) [mynode] (config) #iot transportProfile Sample-Aruba-Telemetry-Https
```
To create an IoT transport profile for Aruba telemetry Websocket, run the following CLI command:
```
(host) [mynode] (config) #iot transportProfile Sample-Aruba-Telemetry-Websocket
```
To apply the IoT transport profile for tracking Meridian beacons, run the following CLI command:
```
(host) [mynode] (config) #iot useTransportProfile Sample-Meridian-Asset-Tracking
```

Creating an IoT Radio Profile

The following procedure describes how to create an IoT radio profile in the WebUI:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > IoT** page.
2. In the **Iot radio profiles** tab, click **+** and configure the following parameters:

Table 256: *IoT Radio Profile Configuration Parameters*

Parameter	Description
Name	Name of the IoT radio profile.
Radio(s)	Type of radio to use. Available options are:

Table 256: *IoT Radio Profile Configuration Parameters*

Parameter	Description
	■ Internal - Use the internal radio in the IoT radio profile. ■ External - Use the external radio in the IoT radio profile. Internal is the default radio type.
Radio mode(s)	Type of the radio mode to use. Available options are: ■ None - Does not use any radio mode. ■ BLE - Use BLE radio mode. ■ Zigbee - Use ZigBee radio mode. ■ BLE & Zigbee - Use both BLE and ZigBee radio modes. None is the default radio mode.
BLE operational mode	BLE operation mode to use. This parameter is available only when **Radio mode** is set to **BLE** or **BLE & Zigbee**. Available options are: ■ Beaconing - Use beaconing in BLE operation mode. ■ Scanning - Use scanning in BLE operation mode. ■ Both - Use both beaconing and scanning modes in BLE operation mode. Beaconing is the default BLE operational mode.
Console	Console mode to use when radio mode is BLE or BLE & Zigbee. This parameter is available only when **Radio mode** is **BLE** or **BLE & Zigbee**. Available options are: ■ Auto - Use console automatically. ■ On - Console is available. ■ Off - Console is not available. Auto is the default Console mode.
Tx power	Tx power in dBM to use. This parameter is available only when **Radio mode** is **BLE** or **BLE & Zigbee**
Zigbee operational mode	Zigbee operation mode to use. This parameter is available only when **Radio mode** is **Zigbee** or **BLE & Zigbee**. The zigbee operation mode is set by default to **coordinator**.
Channel	Channel to use. This parameter is available only when **Radio mode** is **Zigbee** or **BLE & Zigbee**. Available options are: ■ Automatic - Select the channel automatically. ■ Manual - Specify the channel manually. Automatic is the default channel.

3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following command configures an IoT radio profile in the CLI:

```
(host) [mynode] (config) #iot radio-profile <profile-name>
```

The following command displays the list of IoT radio profiles in the CLI:

```
(host) [mynode] #show iot radio-profile
```

The following command displays the details of an IoT transport profile in the CLI:

```
(host) [mynode] #show iot radio-profile <profile-name>
```

Creating a ZigBee Service Profile

The following procedure describes how to create a ZigBee service profile in the WebUI:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > IoT** page.
2. In the **Zigbee service profiles** tab, click **+** and configure the following parameters:

Table 257: *ZigBee Service Profile Configuration Parameters*

Parameter	Description
Name	Name of the ZigBee service profile.
State	State of the ZigBee service profile. Available options are: ■ Enable - ZigBee service profile is enabled. ■ Disable - ZigBee service profile is disabled. Enable is the default state.
PAN ID	PAN ID to use in the ZigBee service profile. Available options are: ■ Automatic - Use an automatic PAN ID. ■ Manual - Manually specify the PAN ID to use. Automatic is the default PAN ID.
Radio(s)	Type of radio to use in the ZigBee service profile. Available options are: ■ Internal - Use the internal radio in the ZigBee service profile. ■ External - Use the external radio in the ZigBee radio profile. ■ Both - Use both external and internal radios in the ZigBee service profile. Internal is the default radio.
Allow device to join	Permit a device to join the network. Available options are: ■ Always - Always allows a device to join the network. ■ On-demand - Allows a device to join the network on based on demand On-demand is the default allow device to join.

3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following command configures a ZigBee service profile in the CLI:

```
(host) [mynode] (config) #zigbee service-profile <profile-name>
```

The following command displays the list of ZigBee service profiles in the CLI:

```
(host) [mynode] #show zigbee service-profile
```

The following command displays the details of a ZigBee service profile in the CLI:

```
(host) [mynode] #show zigbee service-profile <profile-name>
```

Creating a ZigBee Socket Device Profile

The following procedure describes how to create a ZigBee socket device profile in the WebUI:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > IoT** page.
2. In the **Zigbee socket device profiles** tab, click **+**.
3. Under **New Zigbee socket device profile**, enter a name for the ZigBee Socket Device Profile in **Name**.
4. Under the **New Zigbee socket device profile** table, click **+** and configure the following parameters:

Table 258: *ZigBee Socket Device Profile Configuration Parameters*

Parameter	Description
Name	Name of the ZigBee socket device profile.
Direction	Direction of the ZigBee socket device profile. Available options are: ■ Inbound - ZigBee socket device profile is inbound. ■ Outbound - ZigBee socket device profile is outbound. Inbound is the default direction.
Source endpoint	The source endpoint. A source endpoint has to be in the range 1 to 254.
Destination endpoint	The destination endpoint. A destination endpoint has to be in the range 1 to 254.
Profile ID	Profile ID of the ZigBee socket device profile. The Profile ID has to be in the range 0x0000 to 0x7FFF or 0xC000 to 0xFFF
Cluster ID	Cluster ID of the ZigBee socket device profile. The Profile ID has to be in the range 0x0000 to 0x7FFF or 0xC000 to 0xFFF.
APS acknowledge	Allow or disallow AP acknowledge. This parameter is available only when **Direction** is set to **Outbound**. Available options are: ■ Enable - Allow AP acknowledge. ■ Disable - Disallow AP acknowledge.

5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following command configures a ZigBee socket device profile in the CLI:

```
(host) [mynode] (config) #zigbee socket-device-profile <profile-name>
```

The following command displays the list of ZigBee socket device profiles in the CLI:

```
(host) [mynode] #show zigbee socket-device-profile
```

The following command displays the details of a ZigBee socket device profile in the CLI:

```
(host) [mynode] #show zigbee socket-device-profile <profile-name>
```

The following command configures a ZigBee socket inbound profile in the CLI:

```
(host) [mynode] (config) #zigbee socket-inbound-profile <profile-name>
```

The following command displays the list of ZigBee socket inbound profiles in the CLI:

```
(host) [mynode] #show zigbee socket-inbound-profile
```

The following command displays the details of a ZigBee socket inbound profile in the CLI:

```
(host) [mynode] #show zigbee socket-inbound-profile <profile-name>
```

The following command configures a ZigBee socket outbound profile in the CLI:

```
(host) [mynode] (config) #zigbee socket-outbound-profile <profile-name>
```

The following command displays the list of ZigBee socket outbound profiles in the CLI:

```
(host) [mynode] #show zigbee socket-outbound-profile
```

The following command displays the details of ZigBee socket outbound profile in the CLI:

```
(host) [mynode] #show zigbee socket-outbound-profile <profile-name>
```

Creating a BLE Service Profile

The following procedure describes how to create a BLE service profile in the WebUI:

1. In the **Mobility Master** node hierarchy, navigate to the **Configuration > Services** page.
2. In the **IoT** tab, under **BLE service profiles**, click **+** .
3. In the **New BLE service profile** configure the following parameters:

Table 259: *BLE Service Profile Configuration Parameters*

Parameter	Description
Name	Name of the BLE service profile.
Advertising format	Advertisement format of the IoT service. The advertising format is set by default to **ibeacon**.
Radio	Radio to use. Available options are: ■ Internal - Use internal radio in the BLE service profile. ■ External - Use external radio in the BLE service profile. ■ Both - Use both internal and external radios in the BLE service profile. Internal is the default radio.
UUID	UUID to use in the BLE service profile.
Major range	Major range (from and to) of UUID to use in the BLE service profile.
Minor range	Minor range (from and to) of UUID to use in the BLE service profile.

4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following command configures a BLE service profile in the CLI:
```
(host) [mynode] (config) #ble service-profile <profile-name>
```
The following command displays the list of BLE service profiles in the CLI:
```
(host) [mynode] #show ble service-profile
```
The following command displays the details of a BLE service profile in the CLI:
```
(host) [mynode] #show ble service-profile <profile-name>
```

Viewing BLE Status

To view the BLE relay status, run the following CLI command:
```
(host) [mynode] #show ble_relay iot-profile
```
To view the BLE configuration of an AP, run the following CLI command:
```
(host) [mynode] #show ap debug ble-config
```
To view the BLE relay log, run the following CLI command:
```
(host) [mynode] #show ble_relay log
```
To view the BLE relay report, run the following CLI command:
```
(host) [mynode] #show ble_relay report
```
To view the BLE relay jobs, run the following CLI command:
```
(host) [mynode] #show ble_relay jobs
```
To view the BLE relay status for websockets transport, run the following CLI command:
```
(host) [mynode] #show ble_relay disp-attr
```
To view list of BLE relay tag report, run the following CLI command:
```
(host) [mynode] #show ble_relay tag-report
```
To view list of BLE relay WS log, run the following CLI command:

```
(host) [mynode] #show ble_relay ws-log
```

To view third-party devices in the BLE table, run the following CLI command:

```
(host) [mynode] #show ap debug ble-table generic
```


Configuring a ZigBee-USB External Dongle

To configure a ZigBee-USB external dongle, create an IoT radio profile, apply the IoT radio profile, create a service profile, and apply the service profile.

To create an IoT radio profile, run the following CLI command:

```
(host) (config) #iot radio-profile zigbee-usb
(host) (IoT Radio Profile "zigbee-usb") #radio-mode zigbee
(host) (IoT Radio Profile "zigbee-usb") #radio-instance external
(host) (IoT Radio Profile "zigbee-usb") #zigbee-channel auto
(host) (IoT Radio Profile "zigbee-usb") #zigbee-opmode cordinator
```

To check if the IoT radio profile is created, run the following CLI command:

```
(host) #show iot radio-profile zigbee-usb
Name                :zigbee-usb
References          :1
Instance            :external
Mode                :zigbee
BLE Opmode          :scanning beaconing
BLE Console         :
BLE TxPower (dBm)   :0
Zigbee Mode         :coordinator
Zigbee Channel(s)   :auto
```

To apply the IoT radio profile, run the following CLI command:

```
(host) (config) #iot use-radio-profile zigbee-usb
```

To create a service profile, run the following CLI commands:

```
(host) (config) # zigbee service-profile zigbee-usb-test
(host) (Zigbee Service Profile "zigbee-usb-test") # radio-instance external
(host) (Zigbee Service Profile "zigbee-usb-test") # permit-joining on
(host) (Zigbee Service Profile "zigbee-usb-test") # security disable
(host) (Zigbee Service Profile "zigbee-usb-test") # panid auto
```

To apply the service profile, run the following CLI command:

```
(host) (config) #zigbee use-service-profile zigbee-usb-test
```

To check if the service profile is applied, run the following CLI command:

```
(host) #show zigbee service-profile zigbee-usb-test
Name           :zigbee-usb-test
Instance       :external
PAN ID         :auto
Security       :Disable
Permit Joining :on
References     :1
```

If the number of references is more than one, then the service profile is applied.

To check if Zigbee support is onboard or an external USB dongle, run the following CLI commands:

```
(host) #show ap debug zigbee radio-table
Zigbee Radio Table
------------------
IEEE Address             Profile  Onboard  Security  Permit Joining  Channel  Extended
PANID    PANID  RSSI  Num of Clients  Uptime
-----------              -------  -------  --------  --------------  -------  -----------
---    -----  ----  --------------  ------
20:4c:03:ff:fe:8f:46:02  zigbee   No       Disabled  267s            25
204c03fffe8f4602  5a99   --    0               43s
----------------
Total Zigbee radio(s):1
```

To check the firmware of the ZigBee-USB external dongle, run the following CLI commands:

```
(host) #show ap debug ble-config
Radio Configuration
-------------------
Radio Information     NORDIC USB External BLE Zigbee
-----------------     ------------------------------
Radio Profile Type    1
Zigbee Supported      Yes
APB MAC Address       20:4c:03:8f:46:02
APB Radio Profile     0x0002
APB Radio TxPower     0
Bundled BluOS Images  DFU BL(/aruba/bin/beacon_usb_dfu-sdbl_1.3-44.bin) DFU APP
(/aruba/bin/beacon_usb_dfu-app_1.3-44.bin)
```

Chapter 45
AirGroup

AirGroup is a unique enterprise-class capability that leverages zero configuration networking to allow mobile device technologies, like the AirPrint wireless printer service or the AirPlay mirroring service, to communicate over a complex access network topology.



Starting from ArubaOS 8.2.0.0, AirGroup supports profile-based hierarchical configuration. AirGroup runs on Mobility Master, managed devices, or stand-alone controllers. While the Mobility Master-Managed Device deployment model supports centralized mode, distributed mode, or both, the stand-alone controller model supports only the distributed mode.

This section describes the following AirGroup topics:

- Zero Configuration Networking on page 1161
- AirGroup Solution on page 1162
- AirGroup in ArubaOS on page 1162
- AirGroup Services on page 1163
- AirGroup Deployment Models on page 1164
- AirGroup Features on page 1165
- Prerequisites to Enable AirGroup on page 1172
- Configuring AirGroup on page 1179
- Best Practices and Limitations on page 1180
- Troubleshooting and Log Messages on page 1183

Zero Configuration Networking

Zero configuration networking is a technology that enables service discovery, address assignment, and name resolution for desktop computers, mobile devices, and network services. It is designed for flat, single-subnet IP networks such as the home network of a user.

The suite of protocols introduced by Apple® for zero configuration networking over IP is referred to as Bonjour®. Bonjour is supported by most of the Apple product lines including the Mac OS X® operating system, iPhone®, iPod®, iPad®, Apple TV® and AirPort Express®. Bonjour is also included within popular software programs such as Apple iTunes®, Safari, and iPhoto®. Bonjour® can be installed on computers running Microsoft Windows® and is supported by most new network-capable printers.

Bonjour locates devices such as printers, other computers, and the services offered by these devices by using mDNS service records. Bonjour uses the link-scope multicast addresses, so each query or advertisement is limited to a specific VLAN. In large universities and enterprise networks, Bonjour capable devices connect to the network using different VLANs. As a result, an iPad on one enterprise VLAN will not be able to discover the Apple TV that resides on another VLAN. Broadcast and multicast traffic is filtered out of a WLAN network in an effort to reduce network traffic. This inhibits Bonjour (mDNS) services, which rely on multicast traffic.

ArubaOS supports DLNA, a network standard that is derived from UPnP in addition to the mDNS protocol. DLNA uses the Simple Service Discovery Protocol (SSDP) for service discovery on the network. DLNA provides the ability to share digital media between multimedia devices, like Windows and Android,

similar to how mDNS supports Zero Configuration Networking to Apple devices and services. ArubaOS ensures that DLNA seamlessly works with the current mDNS implementation. All the features and policies that are applicable to mDNS are extended to DLNA. This ensures full interoperability between compliant devices.

AirGroup Solution

AirGroup leverages key elements of solution portfolio from Aruba including the ArubaOS and Aruba ClearPass Policy Manager.

AirGroup performs the following functions:

- Enables users to discover network services across IP subnet boundaries in enterprise wireless and wired networks.
- Enables users to access the available AirGroup services such as AirPrint and AirPlay.
- Permits users to access conference room Apple TV during presentations, based on group-based access privileges.
- Provides and maintains seamless connectivity of clients and services across VLANs and SSIDs. It minimizes the mDNS traffic across the wired and wireless network, thereby preserving wired network bandwidth and WLAN airtime.

With AirGroup:

- An AirGroup operator—an end user such as a student can register personal devices. The devices registered by the operator can then automatically be shared with each other.
- Each user can create a user group, such as friends and roommates with whom the user can share the registered devices.
- AirGroup administrators can register and manage shared devices such as printers or conference room Apple TV in an organization. The administrator can grant global access to each device, or limit access based on user name, role, or location.

This chapter provides configuration information for network administrators to enable AirGroup and ClearPass Policy Manager and to register devices with ClearPass Guest.

AirGroup also enables context awareness for services across the network:

- AirGroup is aware of personal devices. An Apple TV in a dorm room, for example, can be associated with the student who owns it.
- AirGroup is aware of shared resources, such as an Apple TV in a meeting room, a printer available to multiple users, or AirPlay in a classroom where a laptop screen is projected on HDTV monitor.
- AirGroup is location aware. For example, an iPad is presented with the closest printer instead of all the printers in a building. If a user in a conference room wants to use an Apple TV to project a MacBook screen on an HDTV monitor, the location-aware AirGroup shows the Apple TV that is closest to that user.

AirGroup in ArubaOS

AirGroup runs as a Loadable Service Module (LSM). The AirGroup application uses the OpenFlow infrastructure to receive the signaling messages from the managed devices. As a LSM, AirGroup can be upgraded to a newer version independent of ArubaOS version.

The following is a list of AirGroup changes:

- Ability to define more than one hop for ap-name based location policy.
- Support for disallowed named VLAN policy for users and servers.
- Extension of support for disallowed VLAN policy for users in addition to servers.
- Extension of support for disallowed role policy for servers in addition to users.
- Enhanced visibility of servers, users, traffic trend, and bandwidth utilization in Dashboard.
- Support for wired users.

AirGroup Services

An administrator may enable or disable individual AirGroup services by using the WebUI or CLI. The following AirGroup services are pre-configured and are available as part of the factory default configuration:

- AirPlay
- AirPrint
- Allowall
- Amazon TV
- DIAL
- DLNA Media
- DLNA Print
- GoogleCast
- iTunes
- RemoteMgmt
- Sharing



By enabling the allowall service, all mDNS or DLNA services, including the default services, are enabled. Aruba does not recommend to enable the allowall service. Enable only the desired services.

The following AirGroup services are enabled by default:

- AirPlay — The AirPlay service allows wireless streaming of music, video, and slide shows from iOS device to Apple TV and other devices that support the AirPlay feature.
- AirPrint — The AirPrint service allows to print from an iPad, iPhone, or iPod Touch directly to any AirPrint compatible printers.
- DIAL — Wi-Fi-enabled streaming devices like Google Chromecast, Roku, Amazon FireTV, and more advertise the Discovery and Launch (DIAL) protocol for clients to search for an available device on a wireless network. Once a device is discovered, the protocol synchronizes information on how to connect to the device. The streaming device connects to a television through an HDMI port to wirelessly streams video and music content to the TV from a smart phone (both Android and iOS), tablet, laptop, or desktop computer devices.

The following AirGroup services are disabled by default:

- DLNA Media — Applications such as Windows Media Player use this service to browse and play media content on a remote device.
- DLNA Print — This service is used by printers which support DLNA.
- GoogleCast — Google Chromecast uses this service to stream video and music content from a smart phone to a TV screen using a wireless network.

- iTunes — iTunes service is used by iTunes Wi-Fi sync and iTunes home-sharing applications across all Apple devices. For best practices, see the Apple iTunes Wi-Fi Synchronization and File Sharing on page 1180.
- RemoteMgmt — Use this service for remote login, remote management, and FTP utilities on Apple devices.
- Sharing — Applications such as disk sharing and file sharing, use the service ID that are part of this service on one or more Apple devices. For best practices, see the Apple iTunes Wi-Fi Synchronization and File Sharing on page 1180.



AirGroup also supports custom and allowall services.

AirGroup Modes

AirGroup supports the following modes:

- Centralized mode
- Distributed mode

Centralized Mode

In centralized mode, the AirGroup service runs on the Mobility Master.



Aruba recommends to use centralized mode in Mobility Master-Managed Device deployment model.

Distributed Mode

In distributed mode, the AirGroup service runs on managed devices where an AirGroup profile is configured. The stand-alone controller deployment model supports only the distributed mode.



Aruba recommends to use distributed mode when the WAN uplink bandwidth between the Mobility Master and managed devices is low and in Branch office deployments.

AirGroup Deployment Models

AirGroup supports following deployment models:

- Mobility Master-Managed Device
- Stand-alone Controller

Mobility Master-Managed Device

Mobility Master is the root of a network hierarchy. A single Mobility Master oversees a number of managed devices that can be co-located or off-campus. In Mobility Master-Managed Device deployment model, AirGroup configuration is allowed on the Mobility Master and Managed Device.

Stand-alone Controller

AirGroup supports domains for stand-alone controllers. This feature, for example, allows iPad users on one stand-alone controller to discover Apple TV available on another stand-alone controller if both stand-alone controllers are part of the same domain. In stand-alone controller deployment model, all AirGroup configuration is allowed only on the managed device.

Deprecated AirGroup Features

The following AirGroup features are deprecated:

- Domain is no longer supported in Mobility Master-Managed Device topology. Domain is supported in stand-alone controller topology.
- Global credits mechanism is removed.
- Active wireless discovery mechanism is removed.
- Location discovery parameter is deprecated.
- mDNS Multicast Response Propagation.

AirGroup Features

This section describes the following AirGroup features:

- Named VLANs on page 1165
- Dashboard on page 1166
- Auto-Association and AirGroup Policy on page 1166
- ClearPass Policy Manager and ClearPass Guest on page 1167
- Group-Based Device Sharing on page 1168
- IPv6 Support on page 1169
- Bluetooth-Based Discovery and AirGroup on page 1169
- DLNA UPnP Support on page 1169
- Domain for Stand-alone Controller on page 1169
- mDNS AP VLAN Aggregation on page 1169
- Scalability on page 1170

Number of Hops

To support location based sharing, AirGroup allows an administrator to define the number of hops or the neighborhood of access points an AirGroup server is shared with. This allows an administrator to deploy and administer AirGroup in large physical places with many access points and clients. An administrator can define the hop count as 1, 2, 3, or no neighborhood.

The hop count policy is available for per server AirGroup policy. Service based auto-associate policy considers single hop RF neighbor APs for visibility of server. The hop count policy is allowed for a maximum of 10 servers when the hop count is 2 or 3. The hop count policy is not available in ClearPass Policy Manager. The multi-hop neighbor table for a server is refreshed every 30 minutes. If an access point is connected or removed in 2 or 3 hop neighborhood of the server, it takes up to 30 minutes for the policy to apply this change.

Named VLANs

Use a named VLAN (which can have a VLAN or a VLAN pool) to define and share relationships in a Mobility Master. For example, a named VLAN "faculty" can have access to AirPlay and AirPrint services

whereas another named VLAN "students" can have access to only the AirPrint service. Only 100 VLAN IDs can be configured per named VLAN.

Named VLANs can be disallowed. The disallowed VLANs can be configured at AirGroup service level or AirGroup global level.

VLAN IDs can be disallowed. The disallowed VLAN IDs can be configured at AirGroup service level and AirGroup global level. The disallow VLAN takes only a single VLAN ID. Any value beyond the range of 1 to 4093 is considered as a named VLAN value.

When global disallow server is configured for a VLAN, then records from any server on this VLAN are not cached.

When global disallow user is configured for a VLAN, then response is not sent from a stand-alone controller for any query from this VLAN.

Roles can be disallowed for users and servers. The disallowed role for users and servers can be configured only at AirGroup service level.

Dashboard

The AirGroup dashboard provides enhanced visibility into AirGroup. The combined view of all AirGroup devices and usage in the network is available under the AirGroup dashboard of every node in the hierarchy. Centralized visibility is available only in Mobility Master-Managed Device topology.

Auto-Association and AirGroup Policy

Auto-association allows AirGroup users to discover nearby AirGroup servers. Auto-association ensures that all the AirGroup users associated to an AP-group, AP-FQLN, or AP and its neighbors discover the AirGroup servers. By default, auto-association is disabled on all AirGroup servers. An administrator can enable auto-association for each AirGroup server separately and configure AP-name, AP-group, or AP-FQLN for auto-association. Auto-association can be enabled for a complete service, which allows all the AirGroup servers who advertise that service to be auto-associated with the configured parameter. If auto-association is enabled, other location-based policy configuration for the AirGroup server on ClearPass Policy Manager or CLI is not honored. Auto-association is applicable only for wireless AirGroup servers.

Auto-associate AP name: When a service is configured with auto-associate AP name. The servers connected to that AP (the AP to which the user is connected) or the RF neighborhood APs are visible to the AirGroup user.

Auto associate AP group: When a service is configured with auto-associate AP group. The servers connected to that AP group (the AP group where the user is connected) or the RF neighborhood APs are visible to the AirGroup user.

Auto associate AP FQLN: When a service is configured with auto-associate AP FQLN. The servers connected to that AP FQLN (the AP FQLN where the user is connected) or the servers in the neighborhood floors are visible to the AirGroup user.

By default, all AirGroup servers are visible to every AirGroup user. AirGroup allows an administrator to configure managed device-based policies for AirGroup servers to limit the visibility of AirGroup servers to destined AirGroup users. To limit the visibility of the AirGroup server to intended AirGroup users, administrator can configure shared user-list, shared role-list, and shared group-list for each AirGroup server.

Administrator can also configure location-based policies for AirGroup devices. For example, administrator can configure if an AirGroup server is visible over a broader area than auto-association configuration. In location-based configuration, administrator can configure AP names, AP groups, and AP FQLNs. Location-based policy configuration limits the visibility of AirGroup server to AirGroup users

who are associated to configured APs, its neighbors, AP-groups, or AP-FQLNs. Administrator can choose whether to consider the neighborhood of the configured AP names.

If an AirGroup policy is configured on ClearPass Policy Manager and node (Mobility Master or managed device), the configuration at the node (Mobility Master or managed device) takes precedence over the configuration at the ClearPass Policy Manager. The stand-alone controller-based policy configuration is persistent for AirGroup.

Consider a setup in which two clients are connected to two APs respectively and both the APs are not RF neighbors.

In scenario 1, when auto-associate AP name is enabled along with enforce registration and AirGroup servers are registered with valid shared parameters, AirGroup users connected to an AP will not be able to discover servers connected to the non-neighboring AP. The users will be only able to discover servers registered to the AP to which they are connected.

When a server is registered with zero AP groups, AirGroup users connected to an AP will not be able to discover servers connected to the non-neighboring AP. The users will be able to discover servers registered to the AP to which they are connected. The user connected to a RF neighbor AP of that of server will be able to discover the server too.

When a server is registered with two AP groups, users connected to AP groups other than the specified ones will not be able to discover the server irrespective of their RF location.

Also, users connected to specified AP groups will be able to discover server irrespective of their RF location.

In scenario 2, when auto-associate AP name is enabled along with enforce registration and AirGroup servers are registered without any parameters, AirGroup users connected to an AP will not be able to discover servers connected to non-neighboring AP. The users will be only able to discover the servers registered to the AP to which they are connected.

A user connected to a RF neighbor AP of that of server will be able to discover the server

In scenario 3, when auto-associate AP name is enabled along with enforce registration and AirGroup servers are registered with invalid shared parameters, AirGroup users connected to an AP will not be able to discover servers connected to non-neighboring ap. The users will not be able to discover the servers registered to the AP to which they are connected as well.

In scenario 4, when auto-associate AP name is enabled along with enforce registration and AirGroup servers are not registered , AirGroup users connected to an AP will not be able to discover servers connected to non-neighboring AP. The users will be able to discover servers registered to the AP to which they are connected. The user connected to a RF neighbor AP of that of server will be able to discover the server too.

In scenario 5, when auto-associate AP name is enabled along with enforce registration and AirGroup servers are registered with AirGroup enabled and policy set to personal with no specified user, the AirGroup users will able to discover servers registered to the RF as well as non-RF AP because the owner of server is equal to the username of user.

In scenario 6, when auto-associate AP name is enabled along with enforce registration and AirGroup servers are registered with AirGroup enabled and policy set to personal with two other users specified (as good as adding no policies), AirGroup users connected to an AP will not be able to discover servers connected to non-neighboring AP. The users will be able to discover servers registered to the AP to which they are connected. The user connected to a RF neighbor AP of that of server will be able to discover the server too. In all cases owner of server is not equal to the username of user.

ClearPass Policy Manager and ClearPass Guest

ClearPass Policy Manager delivers identity and device-based network access control across any wired, wireless, and VPN infrastructure. AirGroup can be deployed with ClearPass Policy Manager

(recommended for large WLANs) or without ClearPass Policy Manager in smaller networks. AirGroup enables context awareness for services across the network and supports a typical customer environment with shared, local, and personal services available to mobile devices.

In centralised mode, RADIUS requests to the CPPM server are sent by the Mobility Master. CoA is sent to the Mobility Master.

In distributed mode , the managed device sends the RADIUS requests. CoA is sent to the managed device.

AirGroup and ClearPass Policy Manager work together to allow users to share personal devices.

- An AirGroup administrator uses ClearPass Policy Manager to authorize end users to register their personal devices.
- An AirGroup operator registers their personal devices (such as an Apple TV) in the ClearPass Guest portal.
- AirGroup enabled Mobility Master sends AirGroup queries to ClearPass Policy Manager for information on the registered devices and associates the access privileges of each device to its allowed services.
- ClearPass Policy Manager sends the CoA to notify the Mobility Master about the registered devices.

For more information on ClearPass Policy Manager, see the *ClearPass Policy Manager User Guide* and *ClearPass Guest Deployment Guide*.

Group-Based Device Sharing

AirGroup supports sharing AirGroup devices such as Apple TV, Printer, and so on to a **User Group** using ClearPass Policy Manager. This is an add-on to the existing device sharing mechanisms such as username, user-role, and location based device sharing using ClearPass Policy Manager. A **User Group** is a logical association of users.

A user can be a part of groups that are defined in an active directory. User group attribute for each user is identified when a user is associated to a wireless network. This is initially identified in authentication module (authentication process). Authentication module sends RADIUS request to RADIUS server as a part of 802.1X authentication and the RADIUS server fetches the user group attribute in the form of VSA from the Active Directory. Subsequently, AirGroup obtains this information from authentication module. This is similar to role of user, however, a user can be a part of more than one groups.

When AirGroup learns about a new device, it interacts with ClearPass Guest to obtain the shared attributes. The shared group(s) attribute is also obtained along with the following attributes:

- Device owner
- Shared location(s)
- Shared user(s)
- Shared role(s)



The group based device sharing feature is supported in ClearPass Policy Manager 6.3 and higher versions.

A user can be a part of maximum 32 user groups. This needs to be defined as comma separated string in Active directory. Each group name can contain a maximum of 63 characters and the entire group name strings cannot exceed 320 characters.

The AirGroup policy engine is enhanced to compare the group membership of the user and shared groups to determine if a user can discover the specific AirGroup server or not.

IPv6 Support

AirGroup supports IPv6 enabled users (for example, iPad) and servers (Apple TV, AirPrint printers). All the AirGroup features are available for both IPv4 and IPv6 clients. On any dual stack client, you must restart the client if the IPv4 interface is disabled.



AirGroup does not support a pure IPv6 deployment in both centralized or distributed modes.

Bluetooth-Based Discovery and AirGroup

Apple devices support Bluetooth-based device discovery mechanism, which allows an Apple device to discover an Apple TV that is within the Bluetooth range.

AirGroup supports only mDNS-based device discovery and does not support Bluetooth-based device discovery mechanism.

DLNA UPnP Support

AirGroup supports DLNA, a network standard that is derived from UPnP in addition to the mDNS protocol.



Cache refresh mechanism is not required for DLNA, as the DLNA devices advertise their service periodically.

Domain for Stand-alone Controller

ArubaOS supports domains for stand-alone controllers. Domain is also supported in Mobility Master topology, when it is running in distributed mode This feature allows, for example, iPad users on one stand-alone controller to discover an Apple TV available on another stand-alone controller if the stand-alone controller is part of the same domain. All stand-alone controllers in a domain communicate with each other.



Configure and use up to 100 AirGroup domains where each AirGroup domain can support up to 100 IP addresses.

AirGroup allows one or more domains to be part of an AirGroup active-domain.

IPv6 is supported in Mobility Master-Managed Device topology but the IPv6 address of Mobility Master is not supported in AirGroup domain. The IPv4 address is supported only in the following scenarios:

- When forming an AirGroup cluster, only IPv4 addresses are supported.
- AirGroup supports IPv4 RADIUS clients only.



The managed device can identify any IPv6 AirGroup servers, only when they proactively advertise their services.

To configure domain for a stand-alone controller, see Configuring AirGroup on page 1179.

mDNS AP VLAN Aggregation

All mDNS or SSDP packets are terminated on a Mobility Master or a stand-alone controller. The AirGroup works as a unicast querier and responder on behalf of mDNS/SSDP devices and eliminates the propagation of multicast mDNS or SSDP traffic in the WLAN.

The mDNS AP VLAN aggregation allows the discovery of wired mDNS or SSDP devices which do not have L2 connectivity with the Mobility Master or a stand-alone controller or which do not trunk on the Mobility Master or a stand-alone controller. An AP, which is in the same VLAN as the wired mDNS or SSDP device which does not trunk on Mobility Master or a stand-alone controller receives and forwards the mDNS or SSDP packets from the wired mDNS or SSDP devices to the Mobility Master or a stand-alone controller. The AP forms a separate split tunnel (0x8000) and aggregates all mDNS or SSDP traffic. Up to 4094 VLANs are supported.

- The split tunnel is formed only when both **AP Multicast Aggregation** (under **AP System Profile**) and **AirGroup** parameters are enabled. If either **AP Multicast Aggregation** or **AirGroup** parameter is disabled, the split tunnel is not formed.
- The **AP Multicast Aggregation** parameter is disabled by default.
- When **AP Multicast Aggregation** parameter is enabled from disabled state, an mDNS or SSDP device discovery packet is sent to the VLAN in which the split tunnel is created if the **AirGroup** parameter is also enabled.
- If an AP is provisioned with an uplink VLAN, then the split tunnel between the AP and the managed device is formed with the uplink VLAN, otherwise the native VLAN is used.
- When the native VLAN is changed, the tunnel is recreated.
- Irrespective of which VLAN (uplink VLAN or native VLAN) is used, the split tunnel is in the same VLAN as the wired mDNS or SSDP devices.
- Configure the VLAN in which the wired mDNS or SSDP device terminates in the managed device. Do not create an SVI or attach a port to the VLAN.

Scalability

AirGroup can scale to support up to 100,000 devices in which up to 17,000 servers can exist.

The following sections describe:

- Scalability Limit in Stand-alone Controller on page 1170
- Memory Utilization on page 1171
- CPU Utilization on page 1171

Scalability Limit in Stand-alone Controller

Table 260 lists the total number of AirGroup servers and users supported by different stand-alone controller models:

Table 260: *AirGroup Server and User Limits in Stand-alone Controller*

Stand-alone Controller Model	Number of AirGroup Servers	Number of AirGroup Users
7240	10000	20000
7220	7000	15000
7210	5000	10000

Stand-alone Controller Model	Number of AirGroup Servers	Number of AirGroup Users
7205	2000	6000
7030	1000	3000
7024	600	1400
7010	500	1500
7005	300	700

The scalability limits are based on:

- Memory Utilization
- CPU Utilization

Memory Utilization

The memory utilization is affected by the number of AirGroup servers and users in an AirGroup cluster. In an AirGroup cluster, the total number of AirGroup servers and users cannot exceed the limit defined by the top-end stand-alone controller. For example, in an AirGroup cluster of one 7005 stand-alone controller and two 7210 stand-alone controllers, the cluster limit is determined as per the scaling limit of the top-end stand-alone controller which is the 7210 stand-alone controller. For the 7005 stand-alone controller in the cluster, the platform limit of the 7005 stand-alone controller is applied.

CPU Utilization

The CPU utilization is measured by the rate at which a stand-alone controller receives mDNS packets per second. The rate of mDNS packets per second in the cluster depends on the number of AirGroup servers, users, and number of applications installed on these devices. When the number of mDNS packets per second exceeds the limit, AirGroup drops the additional packets.

In centralized mode, on Mobility Master, the mDNS packet limit is 1750 pps.

In mixed mode, the mDNS packet limit of 1750 pps applies only for centralized mode. The mDNS packet limit for the platforms depends on controller model (see Table 261).

Table 261 displays the total number of mDNS packets received per second by supported the stand-alone controller or distributed mode:

Table 261: *mDNS Packet Limits in Stand-alone Controller or Distributed Mode*

Stand-alone Controller Model or Distributed Mode	mDNS packets per second (pps)
7280	150
7240	150
7220	90
7210	90
7205	60

Stand-alone Controller Model or Distributed Mode	mDNS packets per second (pps)
7030	75
7024	75
7010	45
7005	45

Use the following command to determine the number of mDNS packets received per second by a managed device:

```
show airgroup internal-state statistics
```



Issue this command multiple times to measure the time difference and the mDNS packet count.

Prerequisites to Enable AirGroup

Before enabling AirGroup, complete the following prerequisites:

- Configure OpenFlow on page 1172
- Bind User VLANs on page 1173
- Enable OpenFlow in User Role and Virtual AP on page 1173
- Configure Management Server Profile on page 1174
- Enable Deep Packet Inspection on page 1175
- Enable Firewall Visibility on page 1175



Complete these prerequisites only on the Mobility Master-Managed Device deployment model.

The following `show` commands can be executed at any specific node path. By default, these `show` commands execute at current node path:

```
show airgroupservice
show airgroup status
show airgroup vlan
show airgroup users
show airgroup servers
show airgroup switches
show airgroup cache entries
show airgroup cppm server-group
show airgroup cppm-server radius statistics
show airgroup cppm-server rfc3576 statistics
```

Configure OpenFlow

By default, OpenFlow is enabled . Enable or Disable OpenFlow on Mobility Master and managed devices. Enable or Disable OpenFlow on the **/mm** node hierarchy.

The following procedure describes how to configure OpenFlow on the Mobility Master:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In **All Profiles**, expand **Controller Profile**.
3. Select **Openflow-controller**.
4. In **Openflow-controller**, select the **ofc-state** check box.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

The following CLI commands configure a OpenFlow on the Mobility Master:

```
(host) [mynode] #cd /mm
(host) [mm] #configure terminal
(host) [mm] (config) #openflow-controller
(host) [mm] (openflow-controller) #openflow-controller-enable
(host) [mm] (openflow-controller) #write memory
```

Bind User VLANs

Bind the user VLANs to the OpenFlow profile on the managed devices. Bind the user VLANs on the **/md** node hierarchy.

The following procedure describes how to bind the user VLANs to the OpenFlow profile on the managed devices:

1. In the **Managed Network** hierarchy, navigate to **Configuration > System > Profiles**.
2. In **All Profiles**, expand **Controller Profile**.
3. Select **Openflow-profile**.
4. In **Openflow-profile**, select the **State** check box.
5. In **controller-ip**, enter the IP address and port number of the Mobility Master.
6. In **bind-vlan**, enter the OpenFlow VLAN to the current list.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

The following CLI commands bind the user VLANs to the OpenFlow profile on the managed devices:

```
(host) [mynode] #cd /md
(host) [md] #configure terminal
(host) [md] (config) #openflow-profile
(host) [md] (Openflow-profile) #openflow-enable
(host) [md] (Openflow-profile) #controller-ip <MM-ip> <port>
(host) [md] (Openflow-profile) #bind-vlan <list of user vlans>
(host) [md] (Openflow-profile) #write memory
```

Enable OpenFlow in User Role and Virtual AP

Enable OpenFlow in the user-role and the virtual AP profile. Enable OpenFlow in the user-role and virtual AP in the **/md** node hierarchy.

The following procedure describes how to enable OpenFlow in the user-role and virtual AP on the managed devices:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Roles & Policies**.
2. In the **Roles** tab, select an existing role.
3. In the **<custom-role>** section below, click **Show Advanced View**.

4. Under **More**, expand **Network**.
5. Select the **Open Flow** check box.
6. Click **Submit**.
7. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
8. In **All Profiles**, expand **Wireless LAN > Virtual AP**.
9. Select **default** profile.



This example uses the default profile.

10. In the default Virtual AP profile, expand **Advanced**.
11. Select the **Openflow Enable** check box.
12. Click **Submit**.
13. Click **Pending Changes**.
14. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

The following CLI commands help to enable OpenFlow in the user-role and virtual AP on the managed devices:

```
(host) [mynode] #cd /md
(host) [mynode] #configure terminal
(host) [md] (config) #user-role <user-role>
(host) [md] (config-submode) #openflow-enable
(host) [md] (config-submode) #!
(host) [md] (config) #wlan virtual-ap <virtual-ap>
(host) [md] (Virtual AP profile "<virtual-ap>") #openflow-enable
(host) [md] (Virtual AP profile "<virtual-ap>") #write memory
```

Configure Management Server Profile

Configure the management server profile to send AMON feeds to the Mobility Master for various statistics. Configure the management server profile in the **/mm** node hierarchy.

The following procedure describes how to configure a management server profile on the Mobility Master:

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > Profiles**.
2. In **All Profiles**, expand **Controller Profile**.
3. Select **Mgmt Config**.
4. In **Mgmt Config profile** window, click **+**.
5. In the **Profile name** field, enter the name of the management server profile.
6. Select the **AirGroup Info** check box.
7. Click **Submit**.
8. In the **Mobility Master** node hierarchy, navigate to **Configuration > System > More > General**.
9. In **MON Receivers**, click **+**.
10. In **New MON Receivers**, enter the following details:
11. In the **Server** field, enter the IP address of the Mobility Master
12. In the **Profile list** drop-down list, select the newly created management server profile.
13. Click **Submit**.
14. Click **Pending Changes**.
15. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

The following CLI commands configure a management server profile on the Mobility Master:

```
(host) [mynode] #cd /mm
(host) [mm] #configure terminal
(host) [mm] (config) #mgmt-server profile <profile-name>
(host) [mm] (Mgmt Config profile "<profile-name>") #airgroup-info enable
(host) [mm] (Mgmt Config profile "<profile-name>") #!
(host) [mm] (config) #mgmt-server primary-server <MM-IP> profile <profile-name>
(host) [mm] (config) #write memory
```

Enable Deep Packet Inspection

Enable Deep Packet Inspection (DPI) on the managed devices. Enable DPI on the **/md** node hierarchy.

The following procedure describes how to enable DPI on the managed devices:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Services > Firewall**.
2. Expand **Global Settings**.
3. Select the **Enable deep packet inspection** check box.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

The following CLI commands help to enable DPI on the managed devices:

```
(host) [mynode] #cd /md
(host) [md] #configure terminal
(host) [md] (config) #firewall
(host) [md] (config-submode) #dpi
(host) [md] (config) #write memory
```

Enable Firewall Visibility

Enable firewall visibility on the managed devices to view the traffic analysis on the Mobility Master dashboard. This is an optional configuration. Enable firewall visibility on the **/md** node hierarchy.

The following procedure describes how to enable firewall visibility on the managed devices:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > Services > Firewall**.
2. Expand **Global Settings**.
3. Select the **Enable firewall visibility** check box.
4. Click **Submit**.
5. Click **Pending Changes**.
6. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

The following CLI commands help to enable firewall visibility on the managed devices:

```
(host) [mynode] #cd /md
(host) [md] #configure terminal
(host) [md] (config) #firewall-visibility
(host) [md] (config) #write memory
```

Configuring AirGroup

AirGroup features are integrated with the WLAN Mobility Master, Managed Device, and Stand-alone Controllers. The Mobility Master also supports optional integration with ClearPass Policy Manager. Trunk all VLANs with wired devices (like printers) on the managed devices.



If your deployment requires ClearPass Policy Manager integration, complete the procedures described in *ClearPass Policy Manager User Guide* and *ClearPass Guest Deployment Guide* before performing the steps described in this section.

Use the following links to configure AirGroup:

- Creating an AirGroup Profile on page 1176
- Creating an AirGroup ClearPass Policy Manager Profile on page 1177
- Creating an AirGroup Domain on page 1178
- Creating an AirGroup Service Profile on page 1178
- Configuring AirGroup on page 1179

The following `show` commands can be executed at any specific node path. By default, these `show` commands execute at current node path:

```
show airgroupservice
show airgroup status
show airgroup vlan
show airgroup users
show airgroup servers
show airgroup switches
show airgroup cache entries
show airgroup cppm server-group
show airgroup cppm-server radius statistics
show airgroup cppm-server rfc3576 statistics
```

Creating an AirGroup Profile

The following procedure describes how to configure an AirGroup profile on a managed device:

1. On the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. Expand **AirGroup** in the **All Profiles** table.
3. Click **Airgroup Profile** under **AirGroup** .
4. In the **AirGroup** profile section click **+**.
5. Type a profile name against **Profile name**.
6. To disallow VLANs, click **+** in **Airgroup disallow vlan** table and in the **Add New** window:
   a. Type VLAN ID or name against **Vlan_id_or_name**.
   b. Select type of AirGroup service in **Airgroup_service** drop-down list.
   c. Select type of users or servers in **Users_servers** drop-down list.
   d. Click **OK**.
7. To disallow roles, click **+** in **Airgroup disallow role** table and in the **Add New** window:
   a. Type role name against **Role_name**.
   b. Select type of AirGroup service in **Airgroup_service** drop-down list.
   c. Select type of users or servers in **Users_servers** drop-down list.
   d. Click **OK**.

8. To auto-associate, click **+** in **Airgroup autoassociate** table and in the **Add New** window:
   a. Select type of AirGroup service in **Airgroup_service** drop-down list.
   b. Select type of auto-association in **Auto_associate** drop-down list.
9. To enforce ClearPass registration, click **Airgroup server enforce registration** check box.
10. Click **Submit**.
11. Click **Pending Changes**.
12. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

Editing an AirGroup Profile

The following procedure describes how to edit an AirGroup profile on a managed device:

1. On the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. Expand **AirGroup** in the **All Profiles** table.
3. Click the desired AirGroup profile under **AirGroup Profile**. In the centralized mode, Click the desired AirGroup profile under **AirGroup**.
4. Edit the following AirGroup profile options:
   - Airgroup disallow vlan
   - Airgroup disallow role
   - Airgroup autoassociate
   - Airgroup server enforce registration
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

Creating an AirGroup ClearPass Policy Manager Profile

AirGroup allows to create and use different ClearPass Policy Manager servers at different nodes.

The following procedure describes how to configure an AirGroup ClearPass Policy Manager profile on a managed device:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. Expand **AirGroup** from the **All Profiles** table.
3. Select **AirGroup ClearPass**.
4. Click **+** in the Airgroup ClearPass window.
5. Type a name for the profile against **Profile name**.
6. Enter a value against **Configure dead time for a down server**.
7. Enter a value against **Configure UDP port to receive RFC 3576 server requests**.
8. Enter a value against **Periodic interval to query ClearPass server**.
9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

You can attach an AirGroup ClearPass Policy Manager profile using the the WebUI.

Attaching an AirGroup ClearPass Policy Manager Profile

The following procedure describes how to attach an AirGroup ClearPass Policy Managerprofile on a managed device:

1. On the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. Expand **AirGroup** in the **All Profiles** table.
3. Expand **AirGroup Profiles** and select the desired profile.
4. Select **AirGroup ClearPass** under the selected profile.
5. Select an AirGroup ClearPass Policy Manager profile from the **Airgroup ClearPass profile** drop-down list.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

Creating an AirGroup Domain

An AirGroup domain is supported only in the distributed mode in all topologies.

The following procedure describes how to configure an AirGroup domain on a managed device:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. Expand **AirGroup** from the **All Profiles** table.
3. Select **AirgGoup Domain**
4. Click **+** in **AirGroup Domain Profile** window.
5. Type an AirGroup domain name against **Profile name**.
6. Type an AirGroup domain description against **Domain description**.
7. To add devices to the domain, click **+** in **Controller ip** table and in the **Add New** window:
   a. Enter the IP address of the devices against **Controller ip**.
   b. Click **OK**.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the required check box and click **Deploy changes**. You can attach an Active-Domain using the WebUI.

Attaching an Active-Domain

The following procedure describes how to attach an active-domain on a managed device:

1. On the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. Expand **AirGroup** in the **All Profiles** table.
3. Expand **AirGroup Profile** and select the desired profile.
4. Select **AirGroup active domain** under the AirGroup profile.
5. Select the AirGroup domain profile in **AirGroup Domain Profile** drop-down list.
6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

Creating an AirGroup Service Profile

The following procedure describes how to create an AirGroup service profile on a managed device:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. Expand **AirGroup** in the **All Profiles** table.
3. Select **Airgroup Service**.
4. Click **+** in **Airgroup Service Profile** section.
5. Type an AirGroup service profile name against **Profile name**.
6. To add an AirGroup service ID, click **+** in **Service Id** table and in the **Add New** window:
7. Enter the AirGroup service ID against **Service Id**.
8. Click **OK**.
9. Type a description for the AirGroup service profile against **Service description**.
10. Click **Submit**.
11. Click **Pending Changes**.
12. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

Attaching an AirGroup Service Profile

The following procedure describes how to attach an AirGroup service on a managed device:

1. In the **Managed Network** node hierarchy, navigate to **Configuration > System > Profiles**.
2. Expand **AirGroup** in the **All Profiles** table.
3. Expand **AirGroup Profile** and select the desired AirGroup profile.
4. Select **AirGroup Service** under the selected AirGroup profile.
5. Click **+** in **Airgroup Service Profile** table.
6. Select the AirGroup service profile from the **Airgroup Service Profile** drop-down list.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

Configuring AirGroup

The following options are available:

- Enabling or Disabling AirGroup on page 1179
- Selecting an AirGroup Mode on page 1179
- Selecting an AirGroup Profile on page 1180

Enabling or Disabling AirGroup

The following procedure describes how to enable or disable AirGroup on a managed device:

1. On the **Managed Network** node hierarchy, navigate to **Configuration > Services > AirGroup**.
2. Click the **AirGroup service** toggle switch to enable or disable.
3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

Selecting an AirGroup Mode

The following procedure describes how to configure AirGroup mode on a managed device:

1. On the **Managed Network** node hierarchy, navigate to **Configuration > Services > AirGroup**.
2. Click the **Centralized** or **Distributed** radio button to select AirGroup to operate in centralized or distributed mode respectively.
3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

Selecting an AirGroup Profile

The following procedure describes how to select an AirGroup profile on a managed device:

1. On the **Managed Network** node hierarchy, navigate to **Configuration > Services > AirGroup**.
2. Select the AirGroup profile from the **AirGroup profile** drop-down list.
3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the required check box and click **Deploy changes**.

Best Practices and Limitations

Consider the best practices and limitations listed in this section before deploying AirGroup. The recommendations that are not specific to a deployment model, apply to both Mobility Master-Managed Device and stand-alone controller deployment model.

- Apple iTunes Wi-Fi Synchronization and File Sharing on page 1180
- Firewall Configuration on page 1181
- Recommended Ports on page 1181
- AirGroup Services for Large Deployments on page 1182
- General AirGroup Limitations on page 1182



Aruba recommends to not enable the allowall service and only enable the desired services.

Aruba recommends to enable autoassociate ap-name in scaled EDU deployments for optimum performance.

Aruba recommends to disallow all roles and VLANS which are not intended to use AirGroup.

Configuring server-based policy is deprecated, instead use ClearPass Policy Manager-based policy.

Apple iTunes Wi-Fi Synchronization and File Sharing

When a managed device receives mDNS response for a service, it caches such records and does not propagate to other users. But for services like iTunes Wi-Fi synchronization and File Sharing to work seamlessly, such mDNS responses must be propagated to other users on the managed device even if they do not query for it.

To ensure that applications such as iTunes Wi-Fi synchronization and File Sharing work seamlessly, ArubaOS selectively forwards these mDNS responses to AirGroup users, based on the user-name

ClearPass Policy Manager policy of the AirGroup server. Hence, for a customer to use these services, it is necessary to configure user-name based ClearPass Policy Manager policies for the AirGroup devices.

Firewall Configuration

The following firewall configuration settings are recommended:

- Disable Inter-User Firewall Settings on page 1181
- ValidUser ACL Configuration on page 1181
- Allow GRE and UDP 5353 on page 1181

Disable Inter-User Firewall Settings

Some firewall settings can prevent the untrusted clients from communicating with each other. When these settings are enabled, an untrusted client such as an iPad may not be able to send its image to an Apple TV on the same managed device.

Use the following commands to disable the virtual AP global firewall options and allow Bonjour services to use AirGroup.

- **no firewall deny-inter-user-bridging**
- **no firewall deny-inter-user-traffic**
- **no ipv6 firewall deny-inter-user-bridging**

ValidUser ACL Configuration

The **ValidUser** ACL must allow mDNS packets with the source IP as a link local address. Do not use a **ValidUser** ACL if the user VLAN interfaces of the AirGroup managed device are not configured with an IP address.

Allow GRE and UDP 5353

mDNS discovery uses the predefined port UDP 5353. If there is a firewall between AirGroup and WLAN, ensure that your firewall policies allow GRE and UDP 5353. DLNA uses the predefined port UDP 1900.

Recommended Ports

The ArubaOS role-based access controls for wireless clients use ACLs to allow or deny user traffic on specific ports. Even though mDNS discovery uses the predefined port UDP 5353, application-specific traffic for services like AirPlay may use dynamically selected port numbers. As a best practice, add or modify ACLs to allow traffic on the ports as described in Table 262 and Table 263.



AirPlay operates using dynamic ports, however, printing protocols like AirPrint use fixed ports.

The following sections describe:

- Ports for AirPlay Service on page 1181
- Ports for AirPrint Service on page 1182

Ports for AirPlay Service

Enable the following ports for the AirPlay services.

Table 262: *Ports for AirPlay Service*

Protocol	Ports
TCP	5000 7000 7100 8612 49152-65535
UDP	7010 7011 8612 49152-65535

Ports for AirPrint Service

Enable the following ports for AirPrint services.

Table 263: *Ports for AirPrint Service*

Protocol	Print Service	Port
TCP	Datastream	9100
TCP	IPP	631
TCP	HTTP	80
TCP	Scanner	9500
TCP	HTTP-ALT	8080

AirGroup Services for Large Deployments

Large deployments with many wireless and wired users often support a large number of advertised Bonjour services, which can consume a significant amount of system resources. For large scale deployments, enable the **AirPlay** and **AirPrint** services, disable the **allowall** service, and then block all other Bonjour services.

General AirGroup Limitations

The AirGroup feature has the following limitations:

- AirGroup is supported only in tunnel and decrypt-tunnel forwarding modes.
- If you use ClearPass Policy Manager to define AirGroup users, the shared user and role lists, and location attributes cannot exceed 1000 characters.
- The RTSP protocol does not support AirPlay on an Apple TV receiver if you enable NAT on the user VLAN interface.
- The location-based access feature only supports AP FQLNs configured in the format **<ap name>.floor <number>.<building>.<campus>**. The AP names cannot contain periods.
- DLNA discovery of AirGroup works across VLANs, however, media streaming from Windows Media Server does not work across VLANs. This limitation is because of Digital Rights Management (DRM) support in Windows Media Server, which restricts media sharing across VLANs. Media streaming works only when both client and server are connected to the same VLAN.

- Android devices cannot discover media server while using the native music and video player applications and when they are connected across VLANs. For example, Samsung Tab 3 cannot discover the media server on Samsung Galaxy S4 while using the native music and video player applications. Android devices can discover media server when they are connected in the same VLAN. This restriction is forced by Samsung devices.
- Xbox cannot be added as an extender to the Windows clients using the Windows Media Center application with the AirGroup feature enabled. You need to disable the AirGroup feature to add Xbox as an extender.
- Wireless Clients such as iPad and iPhone running the Sonos application cannot discover Sonos music system with the AirGroup is enabled.

Troubleshooting and Log Messages

The following procedure describes how to prevent potential AirGroup errors:

1. Execute the **show airgroup internal-state statistics** command and ensure that the **Sibyte Messages Sent/Recv** counters increment over a period of time.
2. Enable mDNS logs using the **logging level debugging system process mdns** command, and capture the output of **show log system all** when the issue occurs. Review any obvious error print statements.
3. Save the output of **show airgroup cache entries** and **show airgroup cppm entries** and look for any discrepancies.

roubleshooting ClearPass Guest

ClearPass Guest includes AirGroup-related events in the application log files. You can configure logging levels to provide debugging information.

1. The following procedure describes how to show debugging information in event logs:
2. In ClearPass Guest, go to **Administration > AirGroup Services** and click the **Configure AirGroup Services** command link. The **Configure AirGroup Services** form opens.
3. In the **AirGroup Logging** drop-down list, select either **Debug—log debug information** or **Trace—log all debug information**. When one of these options is selected, debugging information is provided in the events log.
4. Click **Save Configuration**.

For up-to-date information, see the *ClearPass Guest Deployment Guide*.

Troubleshooting ClearPass Policy Manager

Monitoring and reporting services in ClearPass Policy Manager provide insight into system events and performance.

The following procedure describes how to show incoming AirGroup requests from the managed device:

1. In ClearPass Policy Manager, navigate to **Monitoring > Live Monitoring > Access Tracker**. The **Access Tracker** list view opens.
2. Click a row of the event to view details. The **Summary** tab of the **Request Details** view opens. Additional details may be viewed on the **Input**, **Output**, or **Alerts** tabs, or you can click **Show Logs** to view logging details.

For up-to-date information, see the *ClearPass Policy Manager User Guide*.

Log Messages

Display AirGroup logs by issuing the following CLI commands:

- **show log all**
- **show log system all**
- **show log user all**
- **show log user-debug all**

The log debug messages for the mDNS process are not enabled by default. To enable specific logging levels, use the following CLI commands in configuration mode:

To enable high level mDNS debug messages:

```
(host)(config) #logging level debugging system process mdns
```

To enable mDNS packet processing messages:

```
(host)(config) #logging level debugging system process mdns subcat messages
```

To enable mDNS CLI configuration messages:

```
(host)(config) #logging level debugging system process mdns subcat configuration
```

To enable mDNS Auth and ClearPass Policy Manager user messages:

```
(host)(config) #logging level debugging user process mdns
```

Show Commands

Use the following show commands to view AirGroup configuration data and statistics in the managed device:

Viewing AirGroup Flow Table

```
show airgroup flow-table
```

Viewing AirGroup mDNS and DLAN Cache

```
show airgroup cache entries [mdns|dlna]
```

Viewing AirGroup mDNS and DLNA Statistics

```
show airgroup internal-state statistics [mdns|dlna]
```

Viewing AirGroup VLANs

```
(host) #show airgroup vlan
```

Viewing AirGroup Servers

Use the following command to view the AirGroup server (Apple TV, AirPrint Printer, Google ChromeCast, and so on) status in the managed device:

```
show airgroup servers [dlna|mdns] [verbose]
```

Viewing AirGroup Users

```
show airgroup users [mdns|dlna] [verbose]
```

Viewing Service Queries Blocked by AirGroup

This command displays the service ID that was queried but not available in the AirGroup service table.

```
show airgroup blocked-queries [mdns|dlna]
```

Viewing Blocked Services

The **`airgroup service <servicename> disable`** command disables an AirGroup service by blocking the service IDs for that service. When you enable an AirGroup service, service IDs of that service are enabled automatically. To view the list of blocked services, use the following command:

```
show airgroup blocked-service-id [mdns|dlna]
```

Chapter 46
Instant AP VPN Support

ArubaOS is the companion controller release for the Aruba Instant release. This release provides an ability to terminate VPN and GRE tunnels from Instant AP and provide corporate connectivity to the branch Instant AP network. For more details, see the *Aruba Instant User Guide*.

Starting from ArubaOS 8.3.0.0, IAP-VPN is supported on Mobility Controller Virtual Appliance by using default self-signed certificate (Aruba PKI). For Instant AP to establish IPsec connection with Mobility Controller Virtual Appliance, the controller presents a default self-signed certificate which is uploaded on the Instant AP using Activate.



Mobility Masters (Mobility Master Hardware Appliance and Mobility Master Virtual Appliance) do not support any AP termination including Campus APs, Remote APs and IAP-VPN tunnels.

Through Activate, you can push only one default self-signed certificate to Instant AP which can be used to establish IPsec tunnel with Mobility Controller Virtual Appliance.

VPN features are ideal for:

- enterprises with many branches that do not have a dedicated VPN connection to the Head Quarter.
- branch offices that require multiple APs.
- individuals working from home, connecting to the VPN.

This new architecture and form factor seamlessly adds the survivability feature of Instant APs with the VPN connectivity of RAPs — providing corporate connectivity to branches.

This section includes the following topics:

- Overview on page 1186
- VPN Configuration on page 1191
- Viewing Branch Status on page 1192

Overview

This section provides a brief summary of the new features included in ArubaOS to support VPN termination from Instant AP.

Improved DHCP Pool Management

Instant AP allows you to configure the DHCP address assignment for the branches connected to the corporate network through VPN. In distributed DHCP mode, ArubaOS 8.0 allows designated blocks of IP addresses for static IP users by excluding them from the DHCP scope. In addition, it allows creation of scope of any required size, thereby enabling more efficient utilization of IP address across branches. For detailed information on Distributed DHCP for IAP-VPN, see *Aruba Instant User Guide*.

Termination of Instant AP VPN Tunnels

You can configure Instant APs to terminate VPN tunnels on controllers. When configured, the Instant AP cluster creates a tunnel from the Virtual Controller to an Arubacontroller. However, the controller only acts as a VPN end-point and does not configure the IAP. For more information on how to create a VPN tunnel from a Virtual Controller to an Arubacontroller, see *Aruba Instant User Guide*.

Termination of Instant AP GRE Tunnels

Instant APs have the ability to terminate GRE tunnels on controllers. The Instant AP cluster creates a tunnel from the Virtual Controller to the controller in your corporate office. The controller only acts as a GRE end-point and does not configure the IAP. For more information on how to create a GRE tunnel from Virtual Controller to the controller, see the *Aruba Instant Guide*.

L2/L3 Network Mode Support

The Instant AP functioning as a Virtual Controller enables different DHCP pools (various deployment models) in addition to allocating IP subnets to each branch.

Instant APs support the following DHCP configuration modes:

- L2 Switching Mode: In this mode, Instant AP supports distributed L2 and centralized L2 switching modes of connection to the corporate network. When anInstant AP registers with the controller and has an L2 mode DHCP pool configured, the controller automatically adds the GRE or VPN tunnel associated to this Instant AP into the VLAN multicast table. This allows the clients connecting to this L2 mode VLAN to be part of the same L2 domain on controller.
- L3 Routing Mode: In this mode, Instant AP supports L3 routing mode of connection to the corporate network. The Virtual Controller assigns an IP addresses from the configured subnet and forwards traffic to both corporate and non-corporate destinations. The Instant AP handles the routing on the subnet and also adds a route on the controller after the VPN tunnel is set up during the registration of the subnet. When the Instant AP registers with an L3 mode DHCP pool, the controller automatically adds a route to this DHCP subnet enabling routing of traffic from the corporate network to clients on this VLAN in the branch.

Instant APs support the following DHCP configuration modes:

- L2 Switching Mode: In this mode, Instant AP supports distributed L2 and centralized L2 switching modes of connection to the corporate network. When an Instant AP registers with the controller and has an L2 mode DHCP pool configured, the controller automatically adds the GRE or VPN tunnel associated to this Instant AP into the VLAN multicast table. This allows the clients connecting to this L2 mode VLAN to be part of the same L2 domain on controller.
- L3 Routing Mode: In this mode, Instant AP supports L3 routing mode of connection to the corporate network. The Virtual Controller assigns an IP addresses from the configured subnet and forwards traffic to both corporate and non-corporate destinations. The Instant AP handles the routing on the subnet and also adds a route on the controller after the VPN tunnel is set up during the registration of the subnet. When the Instant AP registers with an L3 mode DHCP pool, the controller automatically adds a route to this DHCP subnet enabling routing of traffic from the corporate network to clients on this VLAN in the branch.

Instant AP VPN Scalability Limits

ArubaOS provides enhancements to the scalability limits for the IAP-VPN branches terminating on the controller. The following table provides the IAP-VPN scalability information for various controller platforms:

Table 264: *IAP-VPN Scalability Limits*

Platforms	IAP-VPN Branches (Preferred)	Route Limit	User Limit (L2 Mode)	VLAN Limit
7280	8,192	32,769	16,384	4,094
7240XM	8,192	32,769	16384	4,094
7220	4,096	16,384	16,384	4,094
7210	2,048	8,192	12,228	4,094
7205	1,024	8,192	8,192	2,048
7030	256	8,189	3,582	256
7024	128	4,093	1,792	128
7010	128	4,093	1,792	128
7008	64	4,093	896	128
7005	64	4,093	896	128

- **IAP-VPN Branches (Preferred)**—The number of IAP VPN branches that can be terminated on a given controller platform.
- **Route Limit**—The number of L3 routes supported on the controller.
- **VLAN Limit**—The number of VLANs supported on the controller.
- **User Limit**—For extended VLANs.

Instant AP VPN OSPF Scaling

ArubaOS allows each IAP-VPN to define a separate subnet derived from a corporate intranet pool to allow IAP-VPN devices to work independently. For information on sample topology and configuration, see OSPFv2.

To redistribute IAP-VPN routes into the OSPF process, use the following command :

```
(host) [mynode] (config) #router ospf redistribute rapng-vpn
```

To verify if the redistribution of the IAP-VPN is enabled, use following command:

```
(host) [mynode] #show ip ospf redistribute
Redistribute RAPNG
```

To configure aggregate route for IAP-VPN routes, use the following command:

```
(host) [mynode] (config) # router ospf aggregate-route rapng-vpn
```

To view the aggregated routes for IAP-VPN routes, use the following command:

```
(host) [mynode]#show ip ospf rapng-vpn aggregate-routes
RAPNG VPN aggregate routes
--------------------------
Prefix Mask Contributing routes Cost
------ ---- ------------------ ----
201.201.200.0 255.255.252.0 5 268779624
100.100.2.0 255.255.255.0 1 10
```

To verify the details of configured aggregated route, use the following command:

```
(host) [mynode]# show ip ospf rapng-vpn aggregated-routes <net> <mask>
(host) [mynode]#show ip ospf rapng-vpn aggregate-routes 100.100.2.0 255.255.255.0
Contributing routes of RAPNG VPN aggregate route
------------------------------------------------
Prefix Mask Next-Hop Cost
------ ---- -------- ----
100.100.2.64 255.255.255.224 5.5.0.10 10
```

To view all the redistributed routes:

```
(host) [mynode]#show ip ospf database
OSPF Database Table
-------------------
Area ID    LSA Type      Link ID        Adv Router     Age   Seq#         Checksum
-------    --------      -------        ----------     ---   ----         --------
0.0.0.15   ROUTER        9.9.9.9        9.9.9.9        159   0x80000016   0xee92
0.0.0.15   ROUTER        10.15.148.12   10.15.148.12   166   0x80000016   0x4c0d
0.0.0.15   NETWORK       10.15.148.12   10.15.148.12   167   0x80000001   0x9674
0.0.0.15   NSSA          12.12.2.0      9.9.9.9        29    0x80000003   0x7b54
0.0.0.15   NSSA          12.12.12.0     9.9.9.9        164   0x80000008   0x63a
0.0.0.15   NSSA          12.12.12.32    9.9.9.9        164   0x80000008   0x7b8
0.0.0.15   NSSA          50.40.40.0     9.9.9.9        164   0x80000007   0x8ed4
0.0.0.15   NSSA          51.41.41.128   9.9.9.9        164   0x80000007   0x68f6
0.0.0.15   NSSA          53.43.43.32    9.9.9.9        164   0x80000007   0x2633
0.0.0.15   NSSA          54.44.44.16    9.9.9.9        164   0x80000007   0x353
N/A        AS_EXTERNAL   12.12.2.0      9.9.9.9        29    0x80000003   0x8c06
N/A        AS_EXTERNAL   12.12.12.0     9.9.9.9        169   0x80000001   0x25e4
N/A        AS_EXTERNAL   12.12.12.32    9.9.9.9        169   0x80000001   0x2663
N/A        AS_EXTERNAL   50.40.40.0     9.9.9.9        169   0x80000001   0xab80
N/A        AS_EXTERNAL   51.41.41.128   9.9.9.9        169   0x80000001   0x85a2
N/A        AS_EXTERNAL   53.43.43.32    9.9.9.9        169   0x80000001   0x43de
N/A        AS_EXTERNAL   54.44.44.16    9.9.9.9        169   0x80000001   0x20fe
```

To verify if the redistributed routes are installed or not.

```
(host) [mynode]#show ip route
Codes: C - connected, O - OSPF, R - RIP, S - static
M - mgmt, U - route usable, * - candidate default, V - RAPNG VPN
Gateway of last resort is Imported from DHCP to network 0.0.0.0 at cost 10
Gateway of last resort is Imported from CELL to network 0.0.0.0 at cost 10
Gateway of last resort is Imported from PPPOE to network 0.0.0.0 at cost 10
Gateway of last resort is 10.15.148.254 to network 0.0.0.0 at cost 1
S*    0.0.0.0/0  [1/0] via 10.15.148.254*
V    12.12.2.0/24 [10/0] ipsec map
V    12.12.12.0/25 [10/0] ipsec map
V    12.12.12.32/27 [10/0] ipsec map
V    50.40.40.0/24 [10/0] ipsec map
V    51.41.41.128/25 [10/0] ipsec map
V    53.43.43.32/27 [10/0] ipsec map
V    54.44.44.16/28 [10/0] ipsec map
C    9.9.9.0/24 is directly connected, VLAN9
C    10.15.148.0/24 is directly connected, VLAN1
C    43.43.43.0/24 is directly connected, VLAN132
C    42.42.42.0/24 is directly connected, VLAN123
C    44.44.44.0/24 is directly connected, VLAN125
C    182.82.82.12/32 is an ipsec map 10.15.149.69-182.82.82.12
C    182.82.82.14/32 is an ipsec map 10.17.87.126-182.82.82.14
```

Branch-ID Allocation

For branches deployed in distributed L3 and distributed L2 modes, the master AP in the branch and the controller must agree upon a subnet and IP addresses to be used for DHCP services in the branch. The process or protocol used by the master AP and the controller to determine the subnet and IP addresses

used in a branch is called BID allocation. If the branches are deployed in local-only or centralized-l2-only modes, the BID allocation process is not required.

The BID allocation process performs the following functions:

- Determines the IP addresses used in a branch for distributed L2 mode
- Determines the subnet used in a branch for distributed L3 mode
- Avoids IP address or subnet overlap (that is, avoids IP conflict)
- Ensures that a branch is allocated the same subnet or range of IP addresses irrespective of which AP in the branch assumes the role of the master in an Instant AP cluster

Trusted Branch Validation

The branch registration requests sent by the Instant AP indicates whether the Instant AP is managed by AirWave or Central, and provides information about the Instant AP configuration modes on AirWave (managed or monitored modes). Based on the information received in the registration request and the status of trusted branch configuration, controller allows the Instant AP to register the branch.

Table 265 maps the Instant AP configuration scenario to the branch registration and trusted branch validation process.

Table 265: *Branch Registration and Trusted Branch Validation*

Instant AP Configuration Mode on AirWave or Central	Trusted Branch DB Validation Enabled on Controller	Branch Registration Allowed
Monitor mode	Yes	No
Monitor mode	No	Yes
Managed mode	Yes	Yes
Managed mode	No	Yes
Not managed by AirWave or Central	Yes	No
Not managed by AirWave or Central	No	Yes

To enable trusted branch DB validation:

- Ensure that the Instant APs are running Instant 4.0 or later.
- Ensure that the controller is upgraded to ArubaOS 8.0 or later.



If you have a master-local setup to be upgraded to ArubaOS 8.0, upgrade the master controller first, and then the local controller.

- Ensure that the IAP-VPN branches are configured through the AirWave or Central management interfaces. If the IAP-VPN branches are not managed by AirWave or Central, or if your network has Instant APs running Instant 3.4 or lower release versions, execute the following command to ensure that all branches are in the trusted list:

```
(host) [mynode] (config) #iap trusted-branch-db allow-all
```

or

```
(host) [mynode] (config) #iap trusted-branch-db add mac-address <mac-address>
```

VPN Configuration

The following VPN configuration steps on the controller enable Instant APs to terminate their VPN connection on the controller:

Whitelist DB Configuration

Controller Whitelist DB

You can use the following CLI command to configure the whitelist DB if the controller is acting as the whitelist entry:

```
(host) [mynode] #whitelist-db rap add mac-address 00:11:22:33:44:55 ap-group test
```

The **ap-group** parameter is not used for any configuration, but needs to be configured. The parameter can be any valid string. If an external whitelist is being used, the MAC address of the AP needs to be saved in the RADIUS server as a lower case entry without any delimiter.

External Whitelist DB

The external whitelist functionality enables you to configure the RADIUS server to use an external whitelist for authentication of MAC addresses of RAPs.

If you are using Windows 2003 server, perform the following steps to configure external whitelist on it. There are equivalent steps available for Windows Server 2008 and other RADIUS servers.

1. Add the MAC addresses for all the RAPs in the Active Directory of the RADIUS server:
   a. Open the **Active Directory and Computers** window, add a new user and specify the MAC address (without the colon delimiter) of the Remote AP for the user name and password.
   b. Right-click the user that you have just created and click **Properties**.
   c. In the **Dial-in** tab, select **Allow access** in the **Remote Access Permission** section and click **OK**.
   d. Repeat Step a through Step b for all RAPs.
2. Define the remote access policy in the Internet Authentication Service:
   a. In the **Internet Authentication Service** window, select **Remote Access Policies**.
   b. Launch the wizard to configure a new remote access policy.
   c. Define filters and select **grant remote access permission** in the **Permissions** window.
   d. Right-click the policy that you have just created and select **Properties**.
   e. In the **Settings** tab, select the policy condition, and **Edit Profile...**.
   f. In the **Advanced** tab, select **Vendor Specific**, and click **Add** to add new vendor specific attributes.
   g. Enter new vendor specific attributes and click **OK**.
   h. In the **IP** tab, provide the IP address of the Remote AP and click **OK**.

VPN Local Pool Configuration

The VPN local pool is used to assign an IP Address to the Instant AP after successful XAUTH VPN.

```
(host) [mynode] (config)# ip local pool "rapngpool" <pool_start_address> <pool_end_
address>
```

Role Assignment for the Authenticated Instant APs

Define a role that includes a source NAT rule to allow connections to the RADIUS server and for the Dynamic RADIUS Proxy in the Instant AP to work. This role is assigned to Instant APs after successful authentication.

```
(host) [mynode] (config) #ip access-list session iaprole
(host) [mynode] (config-submode) #any host <radius-server-ip> any src-nat
(host) [mynode] (config-submode) #any any any permit
(host) [mynode] (config-submode) #!
(host) [mynode] (config) #user-role iaprole
(host) [mynode] (config-submode) #session-acl iaprole
```

VPN Profile Configuration

The VPN profile configuration defines the server used to authenticate the Instant AP (internal or an external server) and the role assigned to the Instant AP after successful authentication.

```
(host) [mynode] (config) #aaa authentication vpn default-iap
(host) [mynode] (VPN Authentication Profile "default-iap") #server-group default
(host) [mynode] (VPN Authentication Profile "default-iap") #default-role iaprole
```



The **default role** parameter of the **aaa authentication vpn** command requires **Policy Enforcement Firewall for VPN users (PEFV)** license.

By default, the controller uses the default Instant AP role. If the administrator changes the Instant AP role name when the IAP's status is UP, then the controller or the Instant AP must be rebooted.

For more information on VPN profile configuration, see the *VPN Configuration* chapter of the *Aruba Instant User Guide*.

Viewing Branch Status

To view the details of the branch information connected to the controller, execute the **show iap table** command.

Sample Configuration

This example shows the details of the branches connected to the controller:

```
(host) [mynode] #show iap table long

IAP Branch Table
----------------
Name              VC MAC Address     Status   Inner IP       Assigned Subnet  Assigned
Vlan
----              --------------     ------   --------       ---------------  ---------
----
Tokyo-CB:D3:16    6c:f3:7f:cc:42:f8  DOWN     0.0.0.0

Paris-CB:D3:16    6c:f3:7f:cc:3d:04  UP       10.15.207.140  10.15.206.99/29  2
LA                6c:f3:7f:cc:42:25  UP       10.15.207.111  10.15.206.24/29  2
```

```
Munich           d8:c7:c8:cb:d3:16  DOWN     0.0.0.0

London-c0:e1     6c:f3:7f:c0:e1:b1  UP       10.15.207.120  10.15.206.64/29  2
Instant-CB:D3    6c:f3:7f:cc:42:1e  DOWN     0.0.0.0

Delhi            6c:f3:7f:cc:42:ca  DOWN     0.0.0.0

Singapore        6c:f3:7f:cc:42:cb  UP       10.15.207.122  10.15.206.120/29 2

Key        Bid(Subnet Name)
---        ----------------
b3c65c...
b3c65c...
b3c65c...  2(10.15.205.0-10.15.205.250,5),1(10.15.206.1-10.15.206.252,5)
a2a65c...  0
b3c65c...  7(10.15.205.0-10.15.205.250,5),8(10.15.206.1-10.15.206.252,5)
b3c65c...
b3c65c...  1(10.15.205.0-10.15.205.250,5),2(10.15.206.1-10.15.206.252,5)
b3c65c...  14(10.15.205.0-10.15.205.250,5),15(10.15.206.1-10.15.206.252,5)
```

The output of this command includes the following parameters:

Table 266: *Instant AP Table Parameters*

Parameter	Description
`Name`	Displays the name of the branch.
`VC MAC Address`	Displays the MAC address of the Virtual Controller of the branch.
`Status`	Displays the current status of the branch (UP or DOWN).
`Inner IP`	Displays the internal VPN IP of the branch.
`Assigned Subnet`	Displays the subnet mask assigned to the branch.
`Assigned Vlan`	Displays the VLAN ID assigned to the branch.
`Key`	Displays the key for the branch, which is unique to each branch.
`Bid(Subnet Name)`	Displays the Branch ID of the subnet. ■ In the example above, the controller displays bid-per-subnet-per-branch i.e., for "LA" branch, BID "2" for the ip-range "10.15.205.0-10.15.205.250" with client count per branch "5"). If a branch has multiple subnets, it can have multiple BIDs. ■ Branches that are in **UP** state and do not have a **Bid(Subnet Name)** means that the Instant AP is connected to a controller which did not assign any bid for any subnet. In the above example, "Paris-CB:D3:16" branch is **UP** and does not have a **Bid(Subnet Name)** information. This means that either the Instant AP is connected to a backup controller or connected to a primary controller without any distributed L2 or L3 subnets. For more information on bid-per-subnet-per-branch and distributed L2 and L3 subnets, see the *DHCP Configuration* chapter of the Aruba *Instant Access Point 6.2.1.0-3.3 User Guide*.



Executing the **show iap table** command does not display the **Key** and **Bid(Subnet Name)** parameters.

Backward Compatibility

ArubaOS 8.4.0.0 introduces the Backward Compatibility feature that enables managed devices to receive register requests on the older TCP port 80. This option is beneficial when managed devices and Instant APs have not been upgraded to ArubaOS 8.4.0.0 simultaneously in a network. When only managed devices are upgraded, users must enable this feature so that managed devices do not drop register requests received on the older HTTP port 80, which can result in service disruption.



To seamlessly process the register requests without causing service disruption, ensure that the Instant AP and managed device are both upgraded to the 8.4.0.0 software version respectively. However, it is highly important that you first upgrade the managed device to the 8.4.0.0 software version, enable backward compatibility on the managed device using the CLI, and only then upgrade the Instant AP to the 8.4.0.0 software version.

Also, you must not upgrade the Instant AP to the 8.4.0.0 software version first when the managed device, terminating on the IAP-VPN is running an older software version.

IAP-VPN CGI handlers are moved to a new port TCP port of 8089. Additionally, Aruba VPN-based authentication is enabled for the IAP-VPN CGI handlers.

You can enable backward compatibility for Instant APs running older versions using the CLI on the Mobility Master:

In the CLI

The following CLI command enables backward compatibility for Instant APs:

```
(host) [mynode] (config) #iapvpn-backward-compatible
```

Chapter 47
External Services Interface

The ESI provides an open interface that is used to integrate security solutions that solve interior network problems such as viruses, worms, spyware, and corporate compliance. ESI allows selective redirection of traffic to external service appliances such as anti-virus gateways, content filters, and intrusion detection systems. When “interesting” traffic is detected by these external devices, it can be dropped, logged, modified, or transformed according to the rules of the device. ESI also permits configuration of different server groups—with each group potentially performing a different action on the traffic.

You can configure ESI to do one or more of the following for each group:

- Redirect specified types of traffic to the server
- Perform health checks on each of the servers in the group
- Perform per-session load balancing between the servers in each group
- Provide an interface for the server to return information about the client that can place the client in special roles such as “quarantine"



ESI cannot function or send information across an IPsec tunnel.

ESI also provides the ESI syslog parser, which is a mechanism for interpreting syslog messages from third-party appliances such as anti-virus gateways, content filters, and intrusion detection systems. The ESI syslog parser is a generic syslog parser that accepts syslog messages from external devices, processes them according to user-defined rules, and then takes configurable actions on system users.

Topics in this chapter include:

- Sample ESI Topology on page 1195
- Understanding the ESI Syslog Parser on page 1197
- Configuring ESI on page 1199
- Sample Route-Mode ESI Topology on page 1207
- Sample NAT-mode ESI Topology on page 1213
- Understanding BRE Syntax on page 1217



The ESI feature requires that the Policy Enforcement Firewall Next Generation license is installed on the managed device.

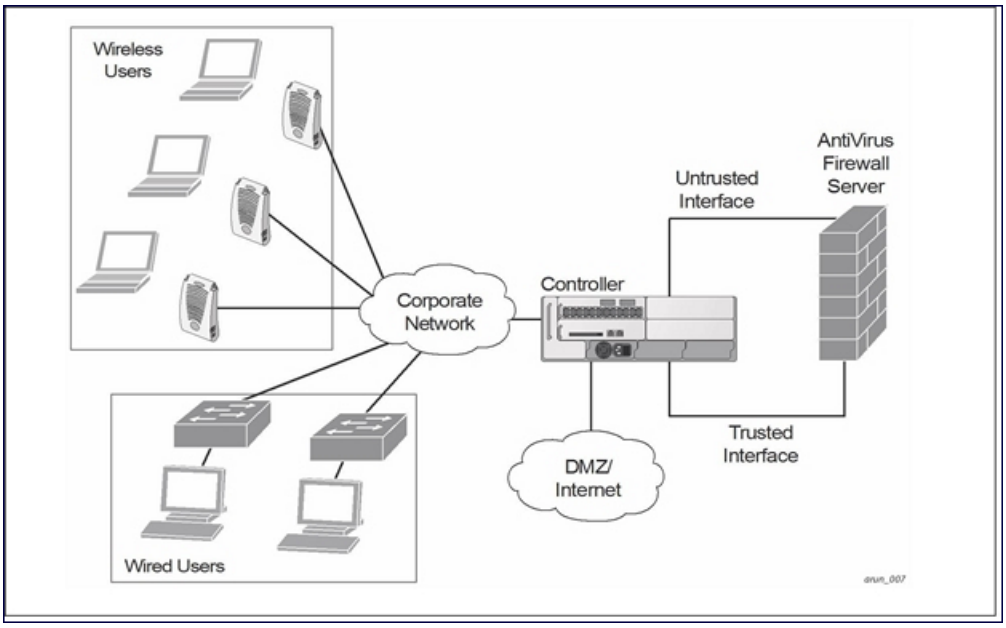
Sample ESI Topology

In the example shown in this section, ESI is used to provide an interface to the antivirus firewall server device for providing virus inspection services. An antivirus firewall server device is one of many different types of services supported in the ESI.



In ArubaOS 3.x, the only antivirus firewall server supported is Fortinet.

Figure 127 *ESI-Fortinet Topology*



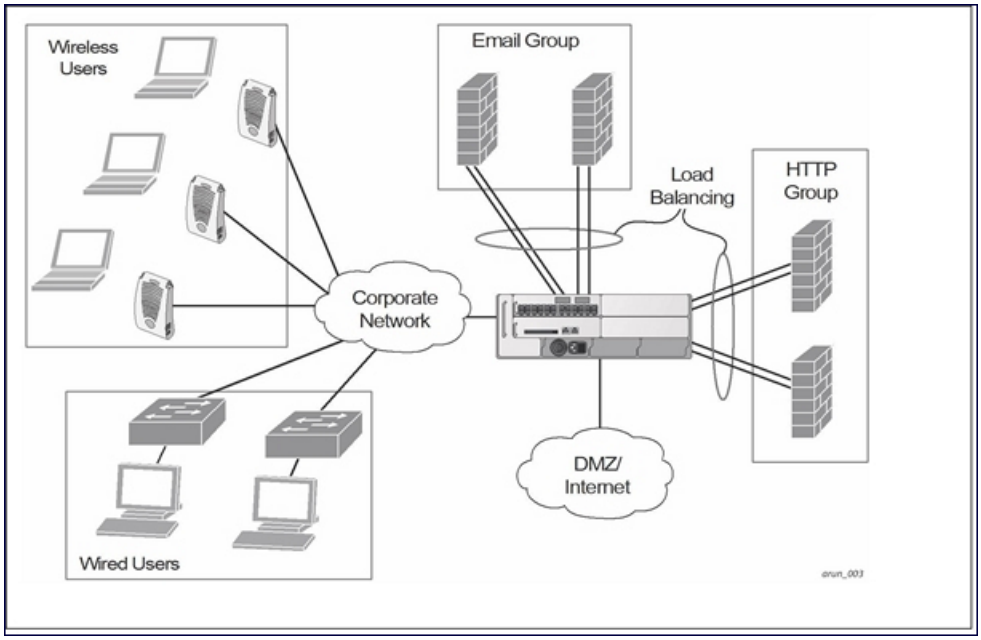
In the ESI–Fortnet topology, the clients connect to access points (both wireless and wired). The wired access points tunnel all traffic back to the managed device over the existing network.

The managed device receives the traffic and redirects relevant traffic (including but not limited to all HTTP or HTTPS and email protocols such as SMTP and POP3) to the antivirus firewall 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 managed device and the antivirus firewall server device. The managed device 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 managed device and the antivirus firewall server device. The managed device forwards all other traffic (for which the antivirus firewall 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 managed device can also be configured to redirect traffic only from clients in a particular role such as “guest” or “non-remediated client” to the antivirus firewall server device. This might be done to reduce the load on the antivirus firewall server device if there is a different mechanism such as the Aruba-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 managed device 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 managed device 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 128 for an example).

Figure 128 *Load Balancing Groups*



Understanding the ESI Syslog Parser

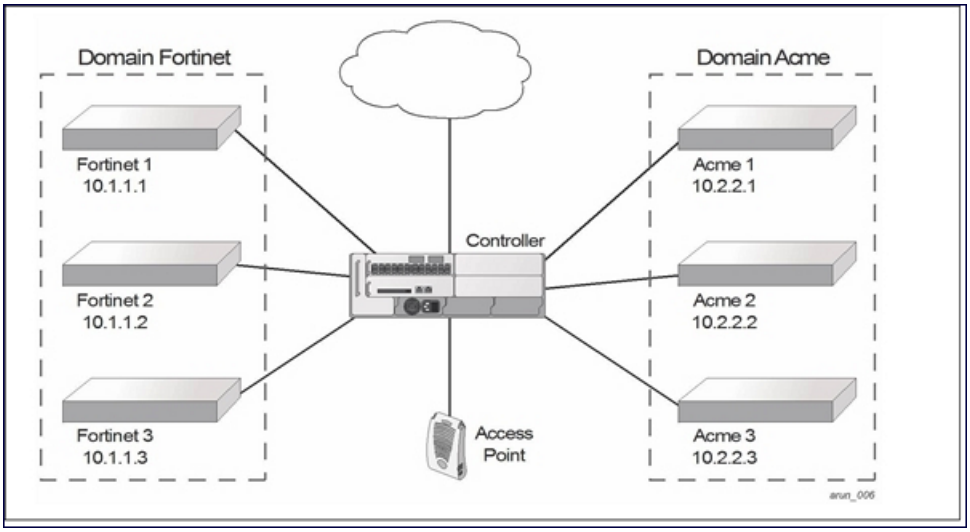
The ESI syslog parser adds a UNIX-style regular expression engine for parsing relevant fields in messages from third-party appliances such as anti-virus gateways, content filters, and intrusion detection systems.

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 129) 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 1199). 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 129 *ESI Parser Domains*



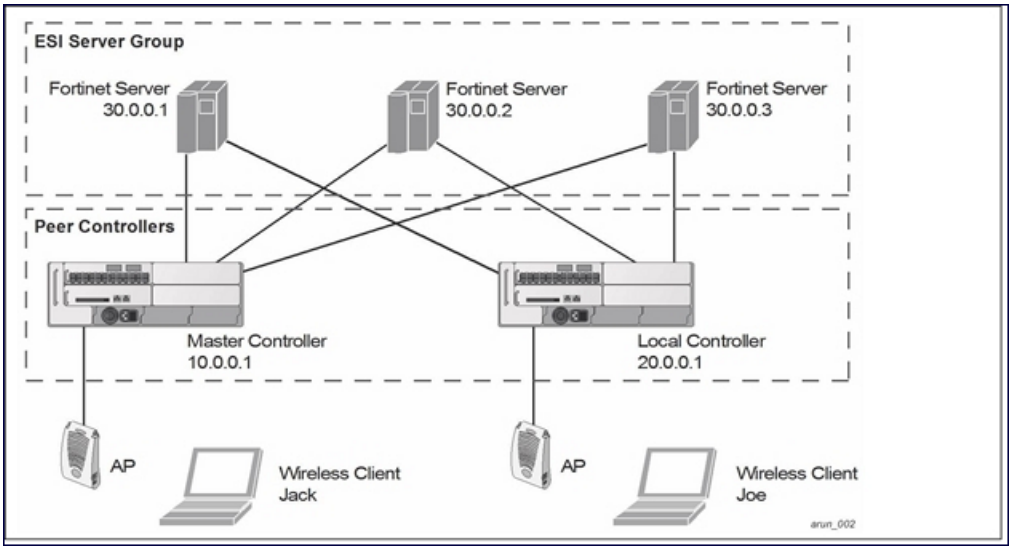
The ESI syslog parser begins with a list of configured IP interfaces which listen for ESI messages. When a syslog message is received, it is checked against the list of defined ESI servers. If a server match is found, the message is then tested against the list of predefined rules.

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 1199). 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 managed device. If the user is not found, the event is meaningless and is ignored. This is the typical situation when a single managed device is connected to a dedicated ESI server.

Peer Managed Devices

As an alternative, consider a topology where multiple managed device share one or more ESI servers.

Figure 130 *ESI Peer managed device*



In this scenario, several managed device (master and local) are defined in the same syslog parser domain to act as *peers*. From the standpoint of the ESI servers, because there is no accurate way of determining from which managed device a given user came. Thus, the event is flooded out to all

managed device defined as peers within this ESI parser domain. The corresponding managed device holding the user entry acts on the event, while other managed device ignore the event.

Syslog Parser Rules

The user creates an ESI rule by using characters and special operators to specify a pattern (regular expression) that uniquely identifies a certain amount of text within a syslog message (Regular expression syntax is described in Understanding BRE Syntax on page 1217. 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 Regex regex() block syntax. This syntax generates two blocks: The first block is the matched expression; the second block contains the value inside the parentheses. For username matching, the focus is on the second block, as it contains the username.

Condition Pattern Matching

The following description uses the Fortigate virus syslog message format as an example to describe condition pattern matching. The Fortigate virus syslog message takes the form:

```
Sep 26 18:30:02 log_id=0100030101 type=virus subtype=infected src=1.2.3.4
```

This message example contains the Fortigate virus log ID number 0100030101 (“log_id=0100030101”), which can be used as the condition—the pattern that uniquely identifies this syslog message.

The parser expression that matches this condition is “log_id=0100030101”. This is a narrow match on the specific log ID number shown in the message, or “log_id=[0–9]{10}[ ]” ,which is a regular expression that matches any Fortigate log entry with a ten-digit log ID followed by a space.

User Pattern Matching

To extract the user identifier in the example Fortigate virus message shown above (“src=1.2.3.4”), use the following expression, “src=(.*)[ ]” to parse the user information contained between the parentheses. The () block specifies where the username will be extracted. Only the first block will be processed.

More examples:

Given a message wherein the username is a MAC address:

```
Sep 26 18:30:02 log_id=0100030101 type=virus subtype=infected mac 00:aa:bb:cc:dd:00
```

The expression “mac[ ](.{17})” will match “mac 00:aa:bb:cc:dd:00” in the example message.

Given a message wherein the username is a user name:

```
Sep 26 18:30:02 log_id=0100030101 type=virus subtype=infected user<johndoe>
```

The expression “user<(.*)>” will match “user<johndoe>” in the example message.

Configuring ESI

You can use the following interfaces to configure and manage ESI and ESI syslog parser behavior:

- The WebUI, which is accessible through a standard Web browser from a remote management console or workstation.
- The CLI, which is accessible from a local console device connected to the serial port on the managed device or through a Telnet or SSH connection from a remote management console or workstation.



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

- The Aruba Management System, which is a suite of applications for monitoring Mobility Master and their related managed devices and APs. Each application provides a Web-based user interface. The Aruba Management System is available as an integrated appliance and as a software application that runs on a dedicated system. See the *ArubaOS User Guide* for more information.

In general, there are three ESI configuration phases on the managed device 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 antivirus firewall.
- The second phase configures the redirection policies instructing the managed device how to redirect the different types of traffic to different server groups.
- The final phase configures the ESI syslog parser domains and the rules that interpret and act on syslog message contents.



The procedures shown in the following sections are general descriptions. Your application might be broader or narrower than this example, but the same general operations apply.

Configuring Health-Check Method, Groups, and Servers

The following procedure describes how to configure health-check method, groups, and server:

To configure the ESI health-check method, servers, and server groups using the WebUI:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > External Services** tab.
2. Expand the **General** accordion.
3. Click **+** in the **Health-Check Configuration** table. The **Create Health-Check** table is displayed (To change an existing profile, click the health check profile name).
4. Provide the following details in the **Create Health-Check** table:
   a. Enter a **Profile name**.
   b. **Frequency (secs)**—Indicates how often the managed device checks to see if the server is up and running. Default: 5 seconds.
   c. **Timeout (secs)**—Indicates the number of seconds the managed device 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 managed device marks the server as being down. Default: 2.
5. Click **Submit** to add a new health check profile.

6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure health-check method, groups, and server:

```
(host) [md] (config) # esi ping profile_name
   frequency seconds
   retry-count count
   timeout seconds
```

Defining the ESI Server

The following procedure describes how to configure an ESI server:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > External Services** tab.
2. Expand the **General** accordion.
3. Click **+** in the **External Servers** table. The **Create Server** section is displayed.
4. Provide the following details in the **Create Server** section:
   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.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

In the CLI

The following CLI commands configure an ESI server:

```
(host) [md] (config)#esi server server_identity
   dport destination_tcp/udp_port
   mode {bridge | nat | route}
   trusted-ip-addr ip-addr [health-check]
   trusted-port <slot/module/port>
   untrusted-ip-addr ip-addr [health-check]
   untrusted-port <slot/module/port>
```

Defining the ESI Server Group

The following procedure describes how to configure an ESI server group:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > External Services** tab.
2. Expand the **General** accordion.
3. Click **+** in the **Server Groups** table. The **Create Server Group** section is displayed.

   (To change an existing group, click the name of the server group you want to edit).
4. Provide the following details in the **Create Server Group** table:
   a. Enter a **Group name**.
   b. In the **Health-Check Profile** drop-down list, select a health check profile.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure an ESI server group:

```
(host) [md] (config)#esi group name
  ping profile_name
  server server_identity
```

Policies and User Role

The following procedure describes how to configure the redirection policies and user role:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Roles** tab.
2. Click **+** in the **Roles** table to create a new user role. The **New Role** section is displayed.

   To change an existing role, click on the role name for the rules to be changed.
3. Enter the **Name** for the role and click **Submit**.
4. To add a policy for the new role, select the name of the role. The **roles > <name of the role>** table is displayed.
5. Click **Show Advanced View** on the table head. The **Policies** tab is now visible.
6. Click **+** in the Policies table. The **Add Policy** section is displayed.

   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 **Create New Policy** option.
   b. Enter the **Policy Name** and click **Submit**.
7. To add a rule to a policy, click the Policy name. The **Role > Policy** table is displayed.
8. Click + to add a new rule. The **New rule for <policy name>** is displayed.
   a. Select the **Rule Type** as **Access Control**. The New rules for the rule type table is displayed.
   b. Select **IPv4** from the **IP Version** drop-down list.
   c. Select **Source**, **Destination**, **Service/app** from their respective drop-down lists.
   d. In the Action drop-down list, select the **Redirect** option.
   e. In the **Redirected to** drop-down list, select **ESI group**.

f. Select **ESI 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).

9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.
12. Refer to Roles and Policies on page 491, for directions on how to apply a policy to a user role.

In the CLI

The following CLI commands configure policies and user role:

```
(host) [md] (config)#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.

(host) [md] (config)#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
```

ESI Syslog Parser Domains and Rules

To view syslog parser domains:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > External Services** view.
2. Click **Syslog Parser Domains**.

   This view lists all the domains by domain name and server IP address, and includes a list of peer managed device.

Adding a new syslog parser domain

The following procedure describes how to add a new syslog parser domain:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > External Services** tab.
2. Expand the **Syslog Parser Domains** accordion.
3. Click **+** in the **Syslog Parser Domains** table. The **New Syslog Parser Domain** is displayed.
4. In the **Domain** text box, type the name of the domain to be added.
5. In the **Servers** field, click **+** in the **IP ADDRESS** table and the **Add Server IP Address** is displayed.
6. In the **Add server IP address** box, enter a valid IP address and click **OK**.



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.

7. In the **Peer controllers** field, click **+** in the **IP ADDRESS** table and the **Add Peer IP Address** is displayed.
8. In the **Add Peer IP address** box, enter a valid IP address and click **OK**.
9. Click **Submit**.
10. Click **Pending Changes**.
11. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands add a new syslog parser domain:

```
(host) [md] (config)#esi parser domain name
  peer peer-ip
  server ipaddr
```

Deleting an existing syslog parser domain

The following procedure describes how to delete an existing syslog parser domain:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > External Services** tab.
2. Expand the **Syslog Parser Domains** accordion.
3. Identify the target parser domain in the list shown in the **Domain** section of the **Syslog Parser Domains** view.
4. Click **Delete** icon on the same row.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command deletes an existing syslog parser domain:

```
(host) [md] (config) #no esi parser domain name
```

Editing an existing syslog parser domain

The following procedure describes how to change an existing syslog parser domain:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > External Services** tab.
2. Expand the **Syslog Parser Domains** accordion.
3. Identify the target parser domain in the list shown in the **Syslog Parser Domains** view .(see To view syslog parser domains: on page 1203).
4. Click on domain name you want to edit and the system displays the edit domain view.



You cannot modify the domain name when editing a parser domain.

5. To delete a server from the selected domain, highlight the server IP address and click **Delete** icon and then click **Submit.**
6. To delete a **Peer controller** server from the selected domain, highlight the **Peer controller** IP address and click **Delete** icon.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes**.

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.

The following CLI commands change an existing syslog parser domain:

```
(host) [md] (config) #esi parser domain name
  no
  peer peer-ip
  server ipaddr
```

The following CLI command shows the syslog parser domain information:

```
(host) [md] #show esi parser domains
```

Managing Syslog Parser Rules

To manage syslog parser rules, click the **Syslog Parser Rules** accordion 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
- + —The actions that can be performed on each rule.

The following procedure describes how to add a new syslog parser rule:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > External Services** tab.
2. Expand the **Syslog Parser Rules** accordion.
3. Click **+** in the **Syslog Parser Rules** table. The **New Syslog Parser Rules** table is displayed.
4. In the **Rule name** text box, type the name of the rule you want to add.
5. Click the **Enable** toggle switch to enable the rule.
6. 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.
7. In the drop-down **Match** list, use the drop-down list to select the match type (ipaddr, mac, or user).
8. 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.
9. 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.
10. In the drop-down **Parser group** list, select one of the configured parser domain names.
11. Click **Submit**.

12. Click **Pending Changes**.
13. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands add a new syslog parser rule:

```
(host) [md] (config) #esi parser rule rule-name
  condition expression
  domain name
  enable
  match {ipaddr expression | mac expression | user expression}
  position position
  set {blacklist | role role}
```

The following procedure describes how 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** icon on the same row.
3. Click **Submit**.
4. Click **Pending Changes**.
5. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI command deletes a syslog parser rule:

```
(host) [md] (config) #no esi parser rule rule-name
```

The following procedure describes how to change an existing syslog parser rule:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > External Services** tab.
2. Expand the **Syslog Parser Rules** accordion.
3. Identify the target parser rule in the list shown in the **Syslog Parser Rules** view.
4. Click on the syslog parser rule name and edit the fields.



You cannot modify the rule name when editing a parser rule.

5. Change the other rule attributes as required:
   a. Click the **Enable** toggle switch to enable the rule.
   b. In the **Condition Pattern** text box, type the regular expression to be used as the condition pattern.
   c. In the drop-down **Match** list, select the match type (ipaddr, mac, or user).
   d. In the **Match Pattern** text box, type the regular expression to be used as the match pattern.
   e. In the drop-down **Set** list, select the set type (blacklist or role).
   f. When you select **role** as the Set type, the system displays a second drop-down list. Click the list to display the possible choices and select the appropriate role value. Validation on the entered value will be based on the Set selection.

g. In the drop-down **Parser Group** list, select one of the configured parser domain names.



At this point, you can test the rule you just edited by using the System Parser Test accordion (accessed from the External Services tab by clicking the Syslog Parser Test accordion, described in Testing a Parser Rule on page 1207.

6. Click **Submit**.
7. Click **Pending Changes**.
8. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands change an existing syslog parser rule:

```
(host) [md] (config) #esi parser rule rule-name
  condition expression
  domain name
  enable
  match {ipaddr expression | mac expression | user expression}
  no
  position position
  set {blacklist | role role}
```

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** tab by clicking the **Syslog Parser Test** tab, which displays the Syslog Parser Rule Test view.

The following procedure describes how to test against a sample syslog message:

1. In the drop-down **Test Type** list, select **Syslog message** as the test source type.
2. In the **Filename** text box, type the syslog message text.
3. 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.

The following procedure describes how to test against a syslog message file:

1. In the drop-down **Test Type** list, select **Syslog file** as the test type.
2. In the **Filename** text box, type the syslog file name.
3. 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.

The following CLI commands test a syslog parser rule:

```
(host) [md] (config) #esi parser rule rule-name
  test {file filename | msg message}
```

The following CLI command shows the syslog parser rule information:

```
(host) [md] #show esi parser rules
```

Monitoring Syslog Parser Statistics

Use the **show esi parser stats** command to monitor syslog parser statistics.

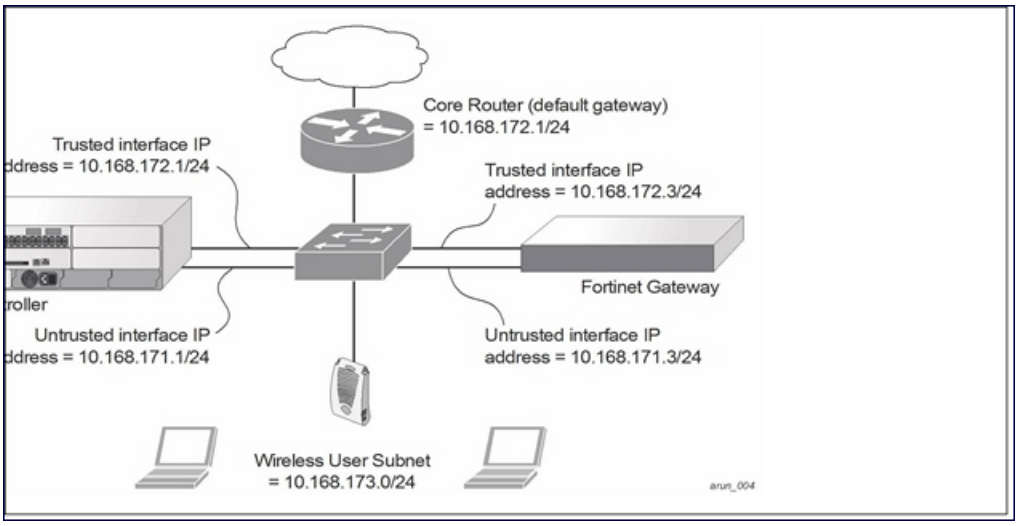
Sample Route-Mode ESI Topology

This section introduces the configuration for a sample route-mode topology using the managed device and Fortinet Anti-Virus gateways. In route mode, the trusted and untrusted interfaces between the managed device and the Fortinet gateways are on different subnets. The following figure shows an example route-mode topology.

NOTE

ESI with Fortinet Anti-Virus gateways is supported only in route mode.

Figure 131 *Example Route-Mode Topology*



In the topology shown, the following configurations are entered on the managed device and Fortinet gateway:

ESI server configuration on the managed device

- Trusted IP address = 10.168.172.3 (syslog source)
- Untrusted IP address = 10.168.171.3
- Mode = route

IP routing configuration on the Fortinet gateway

- Default gateway (core router) = 10.168.172.1
- Static route for wireless user subnet (10.168.173.0/24) through the managed device (10.168.171.2)

Configuring the Example Routed ESI Topology

This section describes how to implement the example routed ESI topology. The description includes the relevant configuration—both the WebUI and the CLI configuration processes are described—required on the managed device to integrate with a antivirus firewall 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 managed device 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 antivirus firewall server devices.
- In the second phase of the configuration task, the user roles are configured with the redirection policies (session ACL definition) instructing the managed device to redirect the different types of traffic to different server groups.
- In the final phase, the ESI parser domains and rules are configured.



The procedures shown in the following sections are based on the requirements in the example routed ESI topology. Your application might be broader or narrower than this example, but the same general operations apply.

Health-Check Method, Groups, and Servers

To configure the ESI health-check method, servers, and server groups, navigate to the **Configuration > Advanced Services > External Services** view on the WebUI.

Defining the Ping Health-Check Method

The following procedure describes how to configure the ping health-check method.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > External Services** tab.
2. Expand the **General** accordion.
3. Click **+** in the **Health-Check Configuration** table. The **Create Health-Check** section is displayed.
4. Provide the following details in the **Create Health-Check** section:
   a. Enter the name **default for the Profile Name.**
   b. **Frequency (secs)**—Enter **5**)
   c. **Timeout (secs)**—Indicates the number of seconds the managed device 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 managed device marks the server as being down. Default: 2 (In this example, enter **3**).
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure the ping health-check method.

```
(host) [md] (config) #esi ping profile_name
   frequency seconds
   retry-count count
   timeout seconds
```

Defining the ESI Server Group

The following procedure describes how to configure an ESI server group.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > External Services** tab.
2. Expand the **General** accordion.
3. Click **+** in the **Server Groups** table. The **Create Server Group** section is displayed.
4. Provide the following details in the **Create Server Group** table:
   a. Enter a **Group name**. Enter **fortinet**.
   b. In the drop-down list, select **default** as the health check profile.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands configure an ESI server group.

```
(host) [md] (config) #esi group name
   ping profile_name
   server server_identity
```

Defining the ESI Server

The following procedure describes how to define an ESI server.

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > External Services** tab.
2. Expand the **General** accordion.
3. Click **+** in the **External Servers** table. The **Create Server** section is displayed.
4. Provide the following details in the **Create Server** section:
   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**.
5. Click **Submit**.
6. Click **Pending Changes**.
7. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands define an ESI server.

```
(host) [md] (config) #esi server server_identity
   dport destination_tcp/udp_port
   mode {bridge | nat | route}
   trusted-ip-addr ip-addr [health-check]
   trusted-port <slot/module/port>
   untrusted-ip-addr ip-addr [health-check]
   untrusted-port <slot/module/port>
```

Redirection Policies and Role

The following procedure describes how to configure the redirection policies and user role.

1. To configure user roles to redirect the required traffic to the server(s), in the **Managed Network** node hierarchy, navigate to the **Configuration > Roles & Policies > Roles** tab.
2. Click **+** to create a new user role.
3. Enter **guest** for Role Name.
4. Click **Submit**.
5. Select **guest** role.
6. Click **Show Advanced View**.
7. Click **+** in **Roles > guest** table.
8. Click **Policies** tab. Click **+** to create a new policy.
9. In the **Add Policy** pop-up, select the **Create a new policy** option. Enter the **Policy Name** as **fortinet** and select **Policy type** as **Session** from the drop-down list.
10. Click **Submit**.
11. Select the **fortinet** policy under the **Roles > guest** table.
12. Click **+** in the **guest Policies > fortinet** table.
13. Select **Access Control** as the **Rule Type** in **New Rule for guest** popup.
14. Enter the following information in the **Roles > fortinet > New forwarding Rule** table.
    - IP version as **IPv4.**
    - Source as **Any**.
    - Destination as **Any**.
    - Service/app as **Protocol** and the Protocol as **svc-http (tcp 80).**
    - Action as **Redirect**.
      - Enter **Redirect to** as **ESI Group**.
      - Enter **Esi group** as **fortinet**.
      - Select **Esi direction** as **Both**. **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.
15. Click **Submit**.
16. Repeat the steps to configure additional rules. This example adds a rule that specifies **any, any, any, permit.**
17. Click **Submit**.
18. Click **Pending Changes**.
19. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands 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.

```
(host) [md] (config) #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.
(host) [md] (config) #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
```

Syslog Parser Domain and Rules

The following sections describe how to configure the syslog parser domain and rules for the route-mode example.

Adding a New Syslog Parser Domain

The following procedure describes how to add a new syslog parser domain for the routed example:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > External services** tab.
2. Expand the **Syslog Parser Domains** accordion and Click **+** in the **Syslog Parser Domains** table. The **New Syslog Parser Domain** is displayed.
3. In the **Domain** text box, type the name of the domain to be added.
4. In the **Servers** field, click **+** in the **IP ADDRESS** table and the **Add Server IP Address** is displayed.
5. In the **Add server IP address** box, enter a valid IP address and click **OK**.



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.

6. In the **Peer Controllers** field, click **+** in the **IP ADDRESS** table and the **Add Peer IP Address** is displayed.
7. In the **Add Peer IP address** box, enter a valid IP address and click **OK**.
8. Click **Submit**.
9. Click **Pending Changes**.
10. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands define a syslog parser domain and the rule to be applied in the route-mode example.

```
(host) [md] (config) #esi parser domain name
   peer peer-ip
   server ipaddr
(host) [md] (config) #esi parser rule rule-name
   condition expression
   domain name
   enable
   match {ipaddr expression | mac expression | user expression }
   position position
   set {blacklist | role role}
```

Adding a New Parser Rule

The following procedure describes how to add a new syslog parser rule for the route-mode example:

1. In the Managed Network node hierarchy, navigate to the **Configuration > Services > External Services** tab.
2. Expand the **Syslog Parser Rules** accordion and click **+** in the **Syslog Parser Rules** table. The **New Syslog Parser Rules** table is displayed.
3. In the **Rule Name** text box, type the name of the rule to be added (in this example, “forti_virus”).
4. Click the **Enable** check box to enable the rule.
5. 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).

6. In the drop-down **Match** list, use the drop-down list to select the match type (in this example, ipaddr).
7. In the **Match Pattern** text box, type the regular expression to be used as the match pattern (in this example, “src=(.*)[ ]”).
8. In the drop-down **Set** list, select the set type (in this example, blacklist).
9. In the drop-down **Parser Group** list, select one of the configured parser domain names (in this example, “forti_domain”).
10. Click **Submit**.
11. Click **Pending Changes**.
12. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Sample NAT-mode ESI Topology

This section describes the configuration for a sample NAT-mode topology using the managed device 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 133.

Figure 132 *Example NAT-Mode Topology*

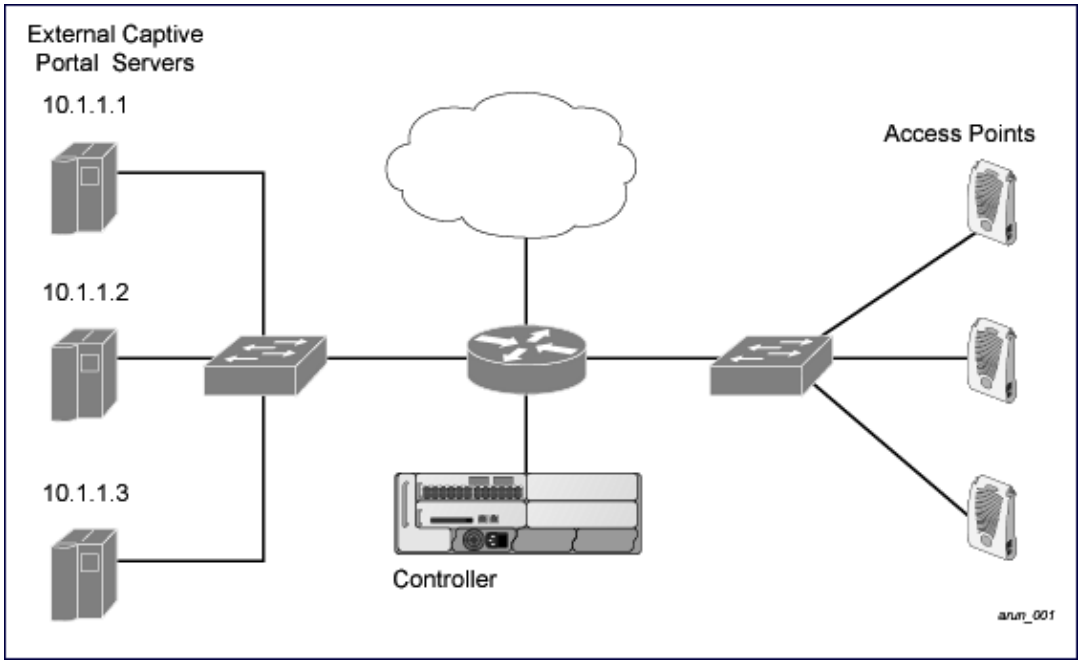
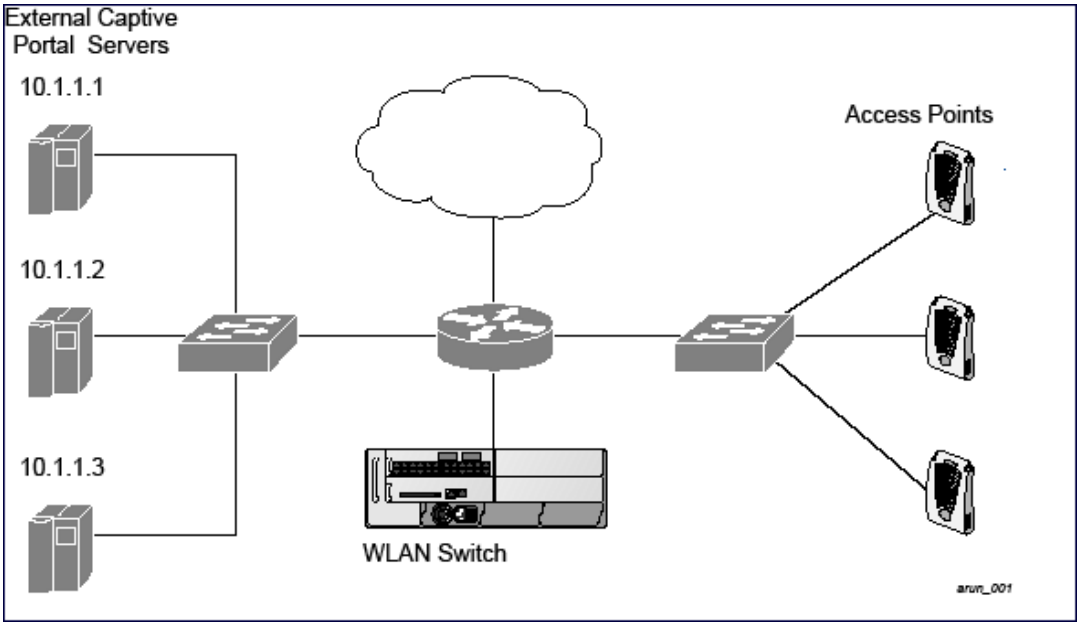


Figure 133



In this example, all HTTP traffic received by the managed device 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 managed device. 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 managed device and external captive-portal servers:

ESI server configuration on the managed device

- 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 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. 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

The following sections describe how to configure NAT-mode ESI example.

Configuring a Health-Check Ping

A health-check ping is associated with an ESI group with servers so that a managed device sends ICMP echo requests to each server in the group and marks the server down if it 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.

The following procedure describes how to configure a health-check ping:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services > External Services** tab.
2. Expand the **General** accordion and Click **+** in the **Health-Check Configuration** table. The **Create Health Check** table is displayed.
3. Provide the following details in the **Create Health-Check** table:
   a. **Profile Name**. This example uses **externalcp_ping**.
   b. **Frequency** seconds. This example uses **30**.
   c. **Retry Count**. This example uses **3**.



If you do not specify a value for a parameter, the WebUI assumes the default value. In this example, the desired timeout value is two seconds; therefore, not specifying the timeout causes the WebUI to use the default value of two seconds.

4. Click **Submit**.

The following CLI commands configure a health-check ping:

```
(host) [md] (config) #esi ping profile_name
```

```
frequency seconds
retry-count count
timeout seconds
```

Configuring the ESI Group

The following procedure describes how to configure a ESI group:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Services** > **External Services** tab.
2. Click **+** in the Server Groups table. The **Create Server Group** table is displayed.
3. Provide the following details in the **Create Server Group** table :
   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**.
4. Click **Submit** when you are finished.

Configuring the ESI Servers

Configuring an ESI server includes configuring server mode to be NAT, configuring the trusted IP address (the server IP address to which packets should be redirected), redirecting to a different port than the original destination port in the packet, and configure an alternate destination port.

1. Click **+** in the **External Servers** table. The **Create Server** table is displayed.
2. Provide the following details in the **Create Server** table:
   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 port**. 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 **Submit** when you are finished.
4. Repeat Step 1 through Step 3 for the remaining external captive portal servers.
5. Click **Submit** to apply the configuration changes.

The following CLI commands configure an ESI server and identify its associated attributes:

```
(host) [md] (config) #esi server server_identity
   dport destination_tcp/udp_port
   mode {bridge | nat | route}
   trusted-ip-addr ip-addr [health-check]
```

The following CLI commands configure an ESI server group, identify its associated ping health-check method, and associate a server with this group:

```
(host) [md] (config) #esi group name
   ping profile_name
   server server_identity
```

Configuring the Redirection Filter

The following procedure describes how to redirect the required traffic to the server.

1. Navigate to **Configuration > Roles & Policies > Roles** tab.
2. Click **+** to create a new user role.
3. Enter **guest** for Role Name.
4. Click **Submit**.
5. Select **guest** role.
6. Click **Show Advanced View**.
7. Click **+** in **Roles >guest** table.
8. Click **Policies** tab. Click **+** to create a new policy.
9. In the **Add Policy** popup, select the **Create a new policy** option. Enter the **Policy Name** as **fortinet** and select **Policy type** as **Session** from the drop-down list.
10. Click **Submit**.
11. Select the **cp_redirect_acl** policy under the **Roles > guest** table.
12. Click **+** in the **guest Policies > cp_redirect_acl** table.
13. Select **Access Control** as the **Rule Type** in **New Rule for guest** popup.
14. Enter the following information in the **Roles > fortinet > New forwarding Rule** table.
    - IP version as **IPv4.**
    - Source as **User**.
    - Destination as **Any**.
    - Service/app as **Service** and the Protocol as **svc-http (tcp 80).**
    - Action as **Redirect**.
      - Enter **Redirect to** as **ESI Group**.
      - Enter **Esi group** as **fortinet**.
      - Select **Esi direction** as **Both**. **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.
15. Click **Submit**.
16. Click **Pending Changes**.
17. In the **Pending Changes** window, select the check box and click **Deploy changes**.

The following CLI commands define the redirection filter for sending traffic to the ESI server.

```
(host) [md] (config) #ip access-list session policy
   user any svc-http redirect esi-group group direction both
```

Understanding BRE Syntax

The ESI syslog parser supports regular expressions created using the BRE syntax described in this section. BRE syntax consists of instructions—character-matching operators (described in Table 267), repetition operators (described in Table 268), or expression anchors (described in Table 269)—used to defined the search or match target.

The sections below provide information on character matching and generally used regular expressions.

Character-Matching Operators

Character-matching operators define what the search will match.

Table 267: *Character-matching operators in regular expressions*

Operator	Description	Sample	Result
.	Match any one character.	grep .ord sample.txt	Matches *ford*, *lord*, *2ord*, etc. in the file sample.txt.
[ ]	Match any one character listed between the brackets	grep [cng]ord sample.txt	Matches only *cord*, *nord*, and *gord*
[^]	Match any one character not listed between the brackets	grep [^cn]ord sample.txt	Matches *lord*, *2ord*, etc., but not *cord* or *nord*
		grep [a-zA-Z]ord sample.txt	Matches *aord*, *bord*, *Aord*, *Bord*, etc.
		grep [^0-9]ord sample.txt	Matches *Aord*, *aord*, etc., but not *2ord*, etc.

Regular Expression Repetition Operators

Repetition operators are *quantifiers* that describe how many times to search for a specified string. Use them in conjunction with the character-matching operators in Table 268 to search for multiple characters.

Table 268: *Regular expression repetition operators*

Operator	Description	Sample	Result
?	Match any character one time if it exists	egrep "?erd" sample text	Matches *berd*, *herd*, etc., *erd*
*	Match declared element multiple times if it exists	egrep "n.*rd" sample.txt	Matches *nerd*, *nrd*, *neard*, etc.
+	Match declared element one or more times	egrep "[n]+erd" sample.txt	Matches *nerd*, *nnerd*, etc., but not *erd*
{n}	Match declared element exactly *n* times	egrep "[a-z]{2}erd" sample.txt	Matches *cherd*, *blerd*, etc., but not *nerd*, *erd*, *buzzerd*, etc.
{n,}	Match declared element at least *n* times	egrep ".{2,}erd" sample.txt	Matches *cherd* and *buzzerd*, but not *nerd*
{n,N}	Match declared element at least *n* times, but not more than *N* times	egrep "n[e]{1,2}rd" sample.txt	Matches *nerd* and *neerd*

Regular Expression Anchors

Anchors describe where to match the pattern, and are a handy tool for searching for common string combinations. Some of the anchor examples use the vi line editor command *:s*, which stands for *substitute*. That command uses the syntax: `s/pattern_to_match/pattern_to_substitute`.

Table 269: *Regular expression anchors*

Operator	Description	Sample	Result
^	Match at the beginning of a line	s/^/blah /	Inserts “blah” at the beginning of the line
$	Match at the end of a line	s/$/ blah/	Inserts “ blah” at the end of the line
\<	Match at the beginning of a word	s/\</blah/	Inserts “blah” at the beginning of the word
		egrep “\<blah” sample.txt	Matches *blahfield*, etc.
\>	Match at the end of a word	s/\>/blah/	Inserts “blah” at the end of the word
		egrep “\>blah” sample.txt	Matches *soupblah*, etc.
\b	Match at the beginning or end of a word	egrep “\bblah” sample.txt	Matches *blahcake* and *countblah*
\B	Match in the middle of a word	egrep “\Bblah” sample.txt	Matches *sublahper*, etc.

References

This implementation is based, in part, on the following resources:

- Lonvick, C., “The BSD syslog Protocol”, RFC 3164, August 2001
- Regular expression reference: http://en.wikipedia.org/wiki/Regular_expression
- Regex syntax summary: http://www.greenend.org.uk/rjk/2002/06/regexp.html
- BRE syntax: http://builder.com.com/5100-6372-1050915.html

This chapter introduces the ArubaOS XML API interface and briefly discusses how you can use simple API calls to perform external user management tasks. Sample scripts are listed at the end of the chapter to help you get started with using the XML API.

Topics in this chapter include:

- Overview on page 1220
- How the ArubaOS XML API Works on page 1220
- Creating an XML Request on page 1221
- XML Response on page 1224
- Sample Scripts on page 1232

Overview

ArubaOS allows you to set up customized external captive portal user management using its native XML API interface. The XML API interface allows you to create and execute user management operations on behalf of the clients or users. You can use the XML API interface to add, delete, authenticate, blacklist, query, or log out a user.

Before you Begin

- XML API requires the PEFNG license.
- Ensure that you have connectivity between your XML API server and the Mobility Master via HTTPS.

How the ArubaOS XML API Works

Typical interaction between your XML API server and Managed Device happens using an HTTPS POST command. A typical communication process using the XML API interface happens as follows:

1. An API command is issued from your server in XML format to Managed Device. The XML request can be composed using a language of your choice using the format described in the Creating an XML Request on page 1221. Sample scripts are available in Python or Bourne Shell, using cURL to generate the HTTPS POST command. See the Sample Scripts on page 1232.
2. The XML request is sent using an HTTPS POST command. The common format of the HTTPS POST is **https://<Managed Device-ip>/auth/command.xml**. See Creating an XML Request on page 1221 for more information.
3. Managed Device processes the XML API request and sends the response to the XML API server. You can use the response and take appropriate action that suits your requirement. The response from Managed Device is returned using predefined formats. See the XML Response on page 1224 for more information.

Configuring an XML Server

The following procedure describes how to configure a RADIUS server:

1. In the **Managed Network** node hierarchy, navigate to the **Configuration > Authentication > Auth Servers** tab.
2. In the **All Servers** table, click **+** to add a new server.
3. Enter a **Name**.
4. Enter the **IP address / hostname** for the new server.
5. Select **XML** from the **Type** drop-down list and click **Submit**.
   The new XML server is displayed in the **All Servers** table.

Creating an XML Request

You can create XML request to add, delete, authenticate, blacklist, query, or logout a user. This section provides XML request formats that you can use for each task.



The XML API functions such as addition, deletion, authentication, blacklisting, querying, and logout have been extended to support IPv6 users in addition to IPv4 users. However, the XML API server must be configured with an IPv4 address for communication with Managed Device.

Adding a User

This XML request uses the **user_add** command to create a new user entry in the Managed Device user table. If the user entry is already present in the user table, the command will modify the entry with the values defined in the XML request, with an exception of IP and MAC address. Session time-out is only applicable to captive portal users.

```
xml=<aruba command="user_add">
   <ipaddr>IP-address_of_the_user</ipaddr>
   <macaddr>MAC-address_of_the_user</macaddr>
   <name>User_Name</name>
   <role>Role_Name</role>
   <session_timeout>Session_timeout</session_timeout>
   <key>Shared_Key</key>
   <authentication>MD5|SHA-1|cleartext</authentication>
   <version>1.0</version>
</aruba>
```

The following options are mandatory when you execute the **user_add** command:

- IP Address
- MAC Address (a valid wireless or wired client on the managed device)
- Key
- Authentication
- Version

Deleting a User



Do not use the **user_delete** command if the intention is to clear the association from the Mobility Master user table. If the client is dual-stack, it re-inherits the authentication state from the IPv6 address. If not dual-stack, the client reverts to the initial role.

This XML request uses the **user_delete** command to delete an existing user from the Mobility Master user table.

```
xml=<aruba command="user_delete">
   <ipaddr>IP-address_of_the_user</ipaddr>
   <macaddr>MAC-address_of_the_user</macaddr>
   <name>User_Name</name>
   <key>Shared_Key</key>
   <authentication>MD5|SHA-1|cleartext</authentication>
   <version>1.0</version>
</aruba>
```

The following options are mandatory when you execute the **user_delete** command:

- IP Address
- Key
- Authentication
- Version



Passing the MAC address or username serves only to perform additional validation against the mandatory IP address. For example, if the IP address is 1.2.3.4 and the MAC address passed is 00:11:22:33:44:55 but the real MAC address is 66:77:88:99:aa:bb then the XML request will fail with response code 1, "**unknown user**" message.

Authenticating a User

This XML request uses the **user_authenticate** command to authenticate against the server group defined in the captive portal profile. This is only applicable to captive portal users.

```
xml=<aruba command="user_authenticate">
   <ipaddr>IP-address_of_the_user</ipaddr>
   <macaddr>MAC-address_of_the_user</macaddr>
   <name>User_Name</name>
   <password>Password_for_the_user</password>
   <key>Shared_Key</key>
   <authentication>MD5|SHA-1|cleartext</authentication>
   <version>1.0</version>
</aruba>
```

The following options are mandatory when you execute the **user_authenticate** command:

- IP Address
- Name
- Password
- Key
- Authentication
- Version



Passing the MAC address serves only to perform additional validation against the mandatory IP address. For example, if the IP address is 1.2.3.4 and the MAC address passed is 00:11:22:33:44:55 but the real MAC address is 66:77:88:99:aa:bb then the XML request will fail with response code 1, "**unknown user**" message.

Blacklisting a User

This XML request uses the **user_blacklist** command to blacklist a user from connecting to your network. This command uses the default blacklist timeout of 3600 seconds. There is no corresponding **clear**

command. You can use the Mobility Master CLI to clear the blacklisted clients. Refer the **show ap blacklist-clients**, **stm remove-blacklist-client**, and **stm purge-blacklist-clients** commands in the *ArubaOS CLI Reference Guide* to clear the blacklisted clients.

```
xml=<aruba command="user_blacklist">
   <ipaddr>IP-address_of_the_user</ipaddr>
   <macaddr>MAC-address_of_the_user</macaddr>
   <name>User_Name</name>
   <key>Shared_Key</key>
   <authentication>MD5|SHA-1|cleartext</authentication>
   <version>1.0</version>
</aruba>
```

The following options are mandatory when you execute the **user_blacklist** command:

- IP Address
- Key
- Authentication
- Version



Passing the MAC address or username serves only to perform additional validation against the mandatory IP address. For example, if the IP address is 1.2.3.4 and the MAC address passed is 00:11:22:33:44:55 but the real MAC address is 66:77:88:99:aa:bb then the XML request will fail with response code 1, "**unknown user**" message.

Querying for User Status

This XML request uses the **user_query** command to get the status and details of a user connected to your network. A dual-stack client can be queried by any of its IPv4 or IPv6 addresses, but only the queried IP address and corresponding values are displayed in the output.

```
xml=<aruba command="user_query">
   <ipaddr>IP-address_of_the_user</ipaddr>
   <macaddr>MAC-address_of_the_user</macaddr>
   <name>User_Name</name>
   <key>Shared_Key</key>
   <authentication>MD5|SHA-1|cleartext</authentication>
   <version>1.0</version>
</aruba>
```

The following options are mandatory when you execute the **user_query** command:

- IP Address
- Key
- Authentication
- Version



Passing the MAC address or username serves only to perform additional validation against the mandatory IP address. For example, if the IP address is 1.2.3.4 and the MAC address passed is 00:11:22:33:44:55 but the real MAC address is 66:77:88:99:aa:bb then the XML request will fail with response code 1, "**unknown user**" message.

Logging Out a User

This XML request uses the **user_logout** command to revert the user to the initial role. This is only applicable to captive portal users. For dual-stack clients, all user-table entries will be reverted to the initial role.

```
xml=<aruba command="user_logout">
   <ipaddr>IP-address_of_the_user</ipaddr>
   <macaddr>MAC-address_of_the_user</macaddr>
   <name>User_Name</name>
   <key>Shared_Key</key>
   <authentication>MD5|SHA-1|cleartext</authentication>
   <version>1.0</version>
</aruba>
```

The following options are mandatory when you execute the **user_logout** command:

- IP Address
- Key
- Authentication
- Version

XML Response

For every successful XML request, Mobility Master will return the processed information as an XML response. There are two types of responses: Default response and Query response.

Default Response Format

The format of a default XML response from Mobility Master is:

```
<aruba>
   <status>Ok | Error</status>
   <code>response_code</code>
   <reason>response_message</reason>
</aruba>
```

In which,

- the status specifies if the XML response succeeds or fails. If the request succeeds, the status tag will contain the **Ok** string. If the request fails, the status tag will contain the **Error** string.
- the code is an integer number that represents the error in the request. This tag is populated only if there is an error in the request.
- the reason is a message that contains descriptive information about the error.

Response Codes

The following response codes are returned if the XML request returns an **Error** string.

Table 270: *XML Response Codes*

Code	Reason message	Description
1	**unknown user** The user specified in the XML request does not exist or is incorrect. If the MAC address or username is specified in the query, Mobility Master restricts the supplied IP address, i.e., the requested user IP address together with any MAC address and username will not be found.	Returned by the **user_authenticate**, **user_delete**, **user_blacklist**, **user_logout**, and **user_query** commands.
2	**unknown role** The specified role in the XML request does not exist in Mobility Master.	Returned by the **user_add** command.
3	**unknown external agent**	This error string is returned due to an unknown source IP (i.e. not configured as an XML server). Or, In case of an **user_add** command, it is likely to be due to the **default-xml-api** AAA profile missing from the AAA authentication wired profile.
4	**authentication failed**	Indicates an authentication failure during **user_authenticate**. This is only applicable to captive portal users.
5	**invalid command** The XML request contains a command not supported by ArubaOS XML API interface.	—
6	**invalid message authentication method** The authentication method specified in the XML request is not supported by the ArubaOS XML API interface.	Returned by commands that contain the authentication method in the XML request.
7	**invalid message digest**	This is due to a mismatch in secret between the XML server and Mobility Master XML API profile. If using non cleartext, this could be an error in the calculation of the hashed secret.
8	**missing message authentication** The authentication method is not specified in the XML request.	Returned by all commands that require the authentication method in the XML request.
9	**missing or invalid version number** The XML request does not contain the version number or the version number is incorrect.	Returned by all commands.
10	**internal error**	—

Code	Reason message	Description
12	**can't use vlan ip**	Indicates the supplied IP matches a VLAN IP on Mobility Master.
13	**invalid ip** The XML request contains invalid IP address of the user or client.	Returned by all commands that required IP address to be specified in the XML request.
14	**can't use switch ip** The XML request contains the Mobility Master IP address instead of the client IP address.	Returned by all commands that required IP address to be specified in the XML request.
15	**missing MAC address** The XML request does not contain the MAC address of the user or client.	Returned by all commands that required MAC address to be specified in the XML request.
16	**unsupported command for this user**	Returned when the requested operation is invalid for the specified user.
17	**socket failed or timed out waiting for operation to complete**	Returned when the status of the requested operation is unavailable; usually signifies a socket communication failure or timeout.

Query Command Response Format

The response of the XML request with the **user_query** command contains detailed information about the status of the user or client.

The **status**, **code** and **reason** values are similar to the default response. The following responses are returned only if the **status** code returns the **Ok** string.

Table 271: *Query Response Code*

Response Code	Description
`status`	Displays the status of the XML response.
`code`	Displays the code as an integer number that represents the error in the request. This tag is populated only if there is an error in the request.
`macaddr`	Displays the MAC address of the client.
`ipaddr`	Displays the IPv4 or IPv6 address of the client.
`name`	Displays the hostname of the user or client.
`role`	Displays the current role of the authenticated client.
`type`	Displays if the client is **wired** or **wireless**.
`vlan`	Displays the VLAN ID of the client.
`location`	Displays the name of the AP to which the client is associated.

Response Code	Description
`age`	Displays the age of the client in Mobility Master. The age is displayed in DD:HH:MM format (Day:Hours:Minutes).
`auth_status`	Displays the authentication status of the client. Available values are: **authenticated** or **unauthenticated**.
`auth_server`	Displays the name of the authentication server used for authenticating the client. This information is available only if the client is authenticated by Mobility Master.
`auth_method`	Displays the authentication mechanism used to authenticate the client. This information is available only if the client is authenticated by Mobility Master.
`essid`	Displays the ESSID to which the client is associated.
`bssid`	Displays the BSSID of the AP to which the client is associated.
`phy_type`	Displays the physical connection type. Available values are: **a**, **b**, **g**, **a-HT**, **g-HT**, and **a-VHT**.
`mobility_state`	Displays the roaming state of the client. Available values are: **Wired (Visitor)**, **Visitor**, **Wired (Away)**, **Away**, **Wired (Foreign VLAN)**, **Foreign VLAN**, **Wired (Remote)**, **Associated (Remote)**, **Wired**, and **Wireless**.
`in_packets`	Displays the total number of incoming packets received by the client.
`in_octets`	Displays the incoming packets (in bytes) received by the client.
`out_packets`	Displays the total number of outgoing packets received by the client.
`out_octets`	Displays the outgoing packets (in bytes) received by the client.

Using the XML API Server

Follow the steps below to use the XML API:

1. Configure an XML API server.
2. Associate the XML API server to an appropriate AAA profile.
3. Configure a user role to direct non-authenticated users to the external captive portal server.
4. Configure captive portal profile and associate that to an initial role (example **logon**).
5. Create an XML request with the appropriate API call.
6. Process XML response appropriately.



The default logon role of a client or user must have captive-portal enabled.

Creating an XML API Request

You can create an XML request with an XML API command and send it to Mobility Master via HTTPS POST. The format of the URL to send the XML request is:

`https://`<Mobility Master-`ip>/auth/command.xml`

- **Mobility Master-ip**: The IP address of Mobility Master that will receive the XML API request
- **command.xml**: The XML request that contains the XML API command.

The format of the XML API request is:

```
xml=<aruba command="<XML API command>">
  <options>Value</options>
  ...
  <options>Value</options>
</aruba>
```

You can specify any of the following commands in the XML request:

Table 272: *XML API Command*

XML API Command	Description
`user_add`	This command creates a new user entry in the Mobility Master user table. If the user entry is already present in the user table, the command will modify the entry with the values defined in the XML request. For an existing user, this command will update any value that is supplied, with an exception of IP and MAC address. Session time-out is only applicable to captive portal users.
`user_delete`	This command deletes an existing user from the Mobility Master user table. **NOTE:** Do not use the **user_delete** command if the intention is to clear the association from the Mobility Master user table. If the client is dual-stack, it re-inherits the authentication state from the IPv6 address. If not dual-stack, the client reverts to the initial role.
`user_authenticate`	This command authenticates against the server group defined in the captive portal profile. This is only applicable to captive portal users.
`user_blacklist`	This command blacklists a user from connecting to your network. This command uses the default blacklist timeout of 3600 seconds. There is no corresponding **clear** command. You can use the Mobility Master CLI to clear the blacklisted clients. Refer the **show ap blacklist-clients**, **stm remove-blacklist-client**, and **stm purge-blacklist-clients** commands in the *ArubaOS CLI Reference Guide* to clear the blacklisted clients.
`user_query`	This command fetches the status and details of a user connected to your network. A dual-stack client can be queried by any of its IPv4 or IPv6 addresses, but only the queried IP address is displayed in the output.
`user_logout`	This command reverts the user to the initial role. This is only applicable to captive portal users. For dual-stack clients, all user-table entries will be reverted to the initial role.

Each XML API command requires certain mandatory options to successfully execute the task. The list of all available options are:

Table 273: *XML API Command Options*

Options	Description	Range / Defaults
`ipaddr`	IP address of the user in IPv4 or IPv6 format.	—

Options	Description	Range / Defaults
`macaddr`	MAC address of the user in aa:bb:cc:dd:ee:ff format.	Enter MAC address with colon.
`user`	Name of the user.	64 character string
`role`	The role to apply to a newly created user, or change of role for an existing user. This option applies to **user_add** and **user_delete** commands only.	64 character string
`password`	The password of the user for authentication.	—
`session_timeout`	Session time-out in seconds. User will be disconnected after this time.	—
`authentication`	Authentication method used to authenticate the message and the sender. You can use any of MD5, SHA-1 or clear text methods of authentication. This option is ignored if shared secret is not configured. It is, however, mandatory if it is configured.	—
`key`	This is the encoded SHA1/MD5 hash of shared secret or plaintext shared secret. This option is ignored if shared secret is not configured on the switch. The actual MD5/SHA-1 hash is 16/20 bytes and consists of binary data. It must be encoded as an ASCII based HEX string before sending. It must be present when the Mobility Master is configured with an xml-api key for the server. Encoded hash length is 32/40 bytes for MD5/SHA-1.	
`version`	The version of the XML API interface available in Mobility Master. This field is mandatory in all XML API requests.	Current version 1.0

Associating the XML API Server to a AAA profile

After you define the XML API server profile associate it to the appropriate AAA profile. If the XML API server is not correctly configured in the appropriate profile, Mobility Master will respond with the **client not authorized** error message. You can add XML API server references to the following AAA profile depending on your requirement:

Creating a AAA Profile

Create a AAA profile for the wireless users and associate the XML API server:

```
(host) [mynode] (config) #aaa profile wirelessusers
(host) ^[mynode] (AAA Profile "wirelessusers") #xml-api-server 10.11.12.13
(host) ^[mynode] (AAA Profile "wirelessusers") #write memory
```

Verify the association of the XML API server to the AAA profile:

```
(host) [mynode] #show aaa profile wirelessusers

AAA Profile "wirelessusers"
---------------------------
Parameter                          Value        Set
---------                          -----        ---
Initial role                       logon
MAC Authentication Profile         N/A
```

```
MAC Authentication Default Role     guest
MAC Authentication Server Group     default
802.1X Authentication Profile       N/A
802.1X Authentication Default Role  guest
802.1X Authentication Server Group  N/A
Download Role from CPPM             Disabled
Set username from dhcp option 12    Disabled
L2 Authentication Fail Through      Disabled
Multiple Server Accounting          Disabled
User idle timeout                   N/A
Max IPv4 for wireless user          2
RADIUS Accounting Server Group      N/A
RADIUS Interim Accounting           Disabled
XML API server                      10.11.12.13
RFC 3576 server                     N/A
User derivation rules               N/A
Wired to Wireless Roaming           Enabled
Device Type Classification          Enabled
Enforce DHCP                        Disabled
PAN Firewall Integration            Disabled
Open SSID radius accounting         Disabled
```

For wireless users, associate the AAA profile to the virtual AP profile:

```
(host) [mynode] (config) #wlan virtual-ap wireless-vap
(host) ^[mynode] (Virtual AP profile "wireless-vap") #aaa-profile wirelessusers
(host) ^[mynode] (Virtual AP profile "wireless-vap") #write memory
```

Verify the association of the AAA profile to the virtual AP profile:

```
(host) [mynode] #show wlan virtual-ap wireless-vap

Virtual AP profile "wireless-vap"
---------------------------------
Parameter                                       Value          Set
---------                                       -----          ---
AAA Profile                     wirelessusers
802.11K Profile                                 default
Hotspot 2.0 Profile                             N/A
Virtual AP enable                               Enabled
VLAN                                            N/A
Forward mode                                    tunnel
SSID Profile                                    default
Allowed band                                    all
Band Steering                                   Disabled
Cellular handoff assist                         Disabled
Openflow Enable                                 Disabled
Steering Mode                                   prefer-5ghz
Dynamic Multicast Optimization (DMO)            Disabled
Dynamic Multicast Optimization (DMO) Threshold  6
Drop Broadcast and Multicast                    Disabled
Convert Broadcast ARP requests to unicast       Enabled
Authentication Failure Blacklist Time           3600 sec
Blacklist Time                                  3600 sec
Deny inter user traffic                         Disabled
Deny time range                                 N/A
DoS Prevention                                  Disabled
HA Discovery on-association                     Enabled
Mobile IP                                       Enabled
Preserve Client VLAN                            Disabled
Remote-AP Operation                             standard
Station Blacklisting                            Enabled
Strict Compliance                               Disabled
VLAN Mobility                                   Disabled
WAN Operation mode                              always
```

```
FDB Update on Assoc                              Disabled
WMM Traffic Management Profile                   N/A
Anyspot profile                                  N/A
```

Create a AAA profile for the wired users and associate the XML API server:

```
(host) [mynode] (config) #aaa profile wiredusers
(host) ^[mynode] (AAA Profile "wiredusers") #xml-api-server 10.11.12.13
(host) ^[mynode] (AAA Profile "wiredusers") #write memory
```

Associate the wired AAA profile to the wired authentication profile:

```
(host) [mynode] (config) #aaa authentication wired
(host) ^[mynode] (Wired Authentication Profile) #profile wiredusers
(host) ^[mynode] (Wired Authentication Profile) #write memory
```

Verify the association of the wired AAA profile to the wired authentication profile:

```
(host) [mynode] #show aaa authentication wired

Wired Authentication Profile
----------------------------
Parameter    Value
---------    -----
AAA Profile  wiredusers
```

For unknown wired users, associate the XML API server to the **default-xml-api** AAA profile.



The **default-xml-api** AAA profile is used only to add or authenticate new users.

Associate the XML API server to the **default-xml-api** AAA profile:

```
(host) [mynode] (config) #aaa profile default-xml-api
(host) ^[mynode] (AAA Profile "default-xml-api") #xml-api-server 10.11.12.13
(host) ^[mynode] (AAA Profile "default-xml-api") #write memory
```

Verify the association of the XML API server to the **default-xml-api** AAA profile:

```
(host) [mynode] #show aaa profile default-xml-api

AAA Profile "default-xml-api" (Predefined (changed))
----------------------------------------------------
Parameter                          Value        Set
---------                          -----        ---
Initial role                       logon
MAC Authentication Profile         N/A
MAC Authentication Default Role    guest
MAC Authentication Server Group    default
802.1X Authentication Profile      N/A
802.1X Authentication Default Role guest
802.1X Authentication Server Group N/A
Download Role from CPPM            Disabled
Set username from dhcp option 12   Disabled
L2 Authentication Fail Through     Disabled
Multiple Server Accounting         Disabled
User idle timeout                  N/A
Max IPv4 for wireless user         2
RADIUS Accounting Server Group     N/A
RADIUS Interim Accounting          Disabled
XML API server                     10.11.12.13
RFC 3576 server                    N/A
User derivation rules              N/A
Wired to Wireless Roaming          Enabled
Device Type Classification         Enabled
Enforce DHCP                       Disabled
PAN Firewall Integration           Disabled
Open SSID radius accounting        Disabled
```

Your Mobility Master is now ready to receive API calls from the XML API server.

Setting up the Captive Portal Profile

Set up a Captive Portal profile with a login page that will redirect users to the external Captive Portal server:

```
(host) [mynode] (config) #aaa authentication captive-portal captive-portal-auth
(host) ^[mynode] (Captive Portal Authentication Profile "captive-portal-auth") #default-role authenticated
(host) (Captive Portal Authentication Profile "captive-portal-auth") #login-page https://10.11.12.13/cgi-bin/login.pl
(host) (Captive Portal Authentication Profile "captive-portal-auth") #switch-in-redirection-url
(host) (Captive Portal Authentication Profile "captive-portal-auth") #write memory
```



The *login-page https://10.11.12.13/cgin-bin/login.pl* is for illustration purposes where the *login.pl* is a Perl script on the external server that handles the external captive portal.

Associating the Captive Portal Profile to an Initial Role

Associate the Captive Portal profile to the logon role:

```
(host) [mynode] (config) #user-role logon
(host) ^[mynode] (config-submode)#captive-portal captive-portal-auth
(host) ^[mynode] (config-submode)#session-acl captiveportal
(host) ^[mynode] (config-submode)#write memory
```

Create an alias for the external Captive Portal server:

```
(host) [mynode] (config) #netdestination xCP
(host) ^[mynode] (config-submode)#host 10.11.12.13
(host) ^[mynode] (config-submode)#write memory
```

You can either create a new ACL or append specific rules to an existing ACLs. Create session ACL for the logon role:

```
(host) [mynode] (config) #ip access-list session captiveportal
(host) ^[mynode] (config-submode)#user alias xCP svc-https permit
(host) ^[mynode] (config-submode)#user alias xCP svc-http permit
(host) ^[mynode] (config-submode)#write memory
```

Monitoring External Captive Portal Usage Statistics

To check the external captive portal authentication statistics, execute the **show aaa xml-api statistics** command. This command displays the number of times an authentication command was executed per client. The command also displays the number of times an authentication event occurred and the number of new authentication events that occurred since the last status check.

```
(host) [mynode] #show aaa xml-api statistics
```

Sample Scripts

You can download the sample scripts from developer.arubanetworks.com . Before downloading the scripts, you must read the following disclaimer.



The sample scripts are examples and provided for illustration purposes only. If you plan to use this script in your environment, ensure that the script meets your IT guidelines. By running this script, you acknowledge that Aruba, a Hewlett Packard Enterprise company is in no way liable for any loss, damage, problems arising from running this script.

The following scripts are available for download:

Python 2.7 Script

- *ArubaXMLDemo.py*: This is a Python 2.7 script. This script demonstrates the basic functionality of the XML API. Using this script, you can send XML requests to add, delete, authenticate, blacklist, query, or log out a user.

Bourne Shell Scripts

- *xml_user_add.sh*: This script adds a user using the **user_add** command.
- *xml_user_del_or_logout.sh*: The **user_delete** part of the script deletes an existing user from the Mobility Master user table. The **user_logout** part of the script reverts an existing user to the initial role in the AAA profile.
- *xml_user_query.sh*: This script fetches the status and details of a user connected in the network using the **user_query** command.



The Bourne Shell scripts work on most Unix, Linux, and Mac operating systems. To run on Windows, you can install Cygwin.

All scripts require cURL to be installed on the XML API server. cURL is an open source command line tool and library for transferring data with URL syntax. You can download cURL from http://curl.haxx.se/download.html.

XML API using Python 2.7

The information covered in the following section is based on running the *ArubaXMLDemo.py* script on a Windows 8.1 64-bit and Python 2.7.

Understanding Request and Response

Mobility Master processes the XML API request and sends the response to the XML API server. The XML response contains the status of the request and a code in case of an error.

Request format: **<script_name> <Mobility Master-ip> <secret_key> <command> [options]**

Understanding XML API Request Parameters

The Table 274 lists all parameters that you can use in a request.

Table 274: *XML API Request Parameters and Descriptions*

Parameter	Description
script_name	The name of the script executable.
Mobility Master-ip	The IP address of Mobility Master that will receive the XML requests.
secret_key	The password used to validate the authentication request from your authentication server. See Using the XML API Server on page 1227 for more information.
command	The XML API command sent to the Mobility Master. You can send one of the following commands per request: ■ **use_add**: Creates a new user entry in the Mobility Master user table. If the user entry is already present in the user table, the command will modify the entry with the

Parameter	Description
	values defined in the XML request. For an existing user, this command will update any value that is supplied, with an exception of IP and MAC address. Session time-out is only applicable to captive portal users. ■ **user_delete**: Deletes an existing user from the Mobility Master user table. **NOTE:** Do not use the **user_delete** command if the intention is to clear the association from the Mobility Master user table. If the client is dual-stack, it re-inherits the authentication state from the IPv6 address. If not dual-stack, the client reverts to the initial role. ■ **user_authenticate**: Authenticates against the server group defined in the captive portal profile. This is only applicable to captive portal users. ■ **user_blacklist**: Blacklists a user from connecting to your network. This command uses the default blacklist timeout of 3600 seconds. There is no corresponding **clear** command. You can use the Mobility Master CLI to clear the blacklisted clients. Refer the **show ap blacklist-clients**, **stm remove-blacklist-client**, and **stm purge-blacklist-clients** commands in the *ArubaOS CLI Reference Guide* to clear the blacklisted clients. ■ **user_query**: Fetches the status and details of a user connected to your network. A dual-stack client can be queried by any of its IPv4 or IPv6 addresses, but only the queried IP address is displayed in the output. ■ **user_logout**: Reverts the user to the initial role. This is only applicable to captive portal users. For dual-stack clients, all user-table entries will be reverted to the initial role.
Options	■ **-i <ip_addr>**: Specify the IP address of the user in IPv4 or IPv6 format. ■ **-m <mac_addr>**: Specify the MAC address of the user in aa:bb:cc:dd:ee:ff format. ■ **-n <name>**: Specify the name of the user. ■ **-p <password>**: Specify the password of the user for authentication. ■ **-r role**: Specify the role to apply to a newly created user, or change of role for an existing user. This option applies to **user_add** and **user_delete** commands only. ■ **-t timeout**: Specifies the session time-out in seconds. User will be disconnected after this time. ■ **-v version**: Specifies the version of the XML API interface available in Mobility Master. This field is mandatory is all requests. Default version is 1.0. ■ **-a method**: Specifies the encryption method to send the secret key. You can specify MD5 or SHA-1 or cleartext as the encryption method. By default, cleartext method is used to send the key. ■ **-s sessid**: Specifies the active session ID.

Understanding an XML API Response

The response message from Mobility Master is sent in an XML format. The default format of the response is:

```
[Message header]
Displays the request parameters and other standard header details.
...
...
...
```

```
<response>
        <status>Status Message</status>
        <code>Code in case of an error</code>
</response>
```

The following section describes few of the XML API requests and responses from Mobility Master.

Adding a User

This XML request uses the **user_add** command to create a new user entry in the Mobility Master user table.

```
C:\Python27>python ArubaXMLDemo.py --switch-ip=192.0.2.1 --secret=aruba123 --
command=user_add --ip=192.0.2.2 --mac=a4:e:60:c3:10:59 --role=logon
```

The command sends the following information in the XML request to Mobility Master:

- **--switch-ip**: IP address of Mobility Master
- **--secret**: Shared secret key (sent as plain text)
- **--command**: XML API command
- **--ip**: IP address of the user
- **--mac**: MAC address of the user
- **--role**: User role

Mobility Master Response

Mobility Master processes using an XML format and sends the following response to the XML API server.

```
Warning: The specified mac address *must* match the user specified by --ip or the command
will fail.

Prepared XML buf
---------------------------------
xml=<aruba command="user_add">
<ipaddr>192.0.2.2</ipaddr>
<macaddr>a4:5e:60:c3:10:59</macaddr>
<role>logon</role>
<authentication>cleartext</authentication>
<key>aruba123</key>
<version>1.0</version>
</aruba>
---------------------------------
Sending XML request to https://192.0.2.1/auth/command.xml
Controller response status: 200
Response XML
-----------------------------------
<aruba>
<status>Ok</status>
<code>0</code>
</aruba>
-----------------------------------
```

Mobility Master

You can view the updated details of the user in the Mobility Master.

```
(host) [mynode] #show user-table

Users
-----
IP           MAC               Name    Role  Age(d:h:m)  Auth  VPN link  AP name
----------   ------------      ------  ----  ----------  ----  --------  -------
192.0.2.2    a4:5e:60:c3:10:59         logon 00:00:00
```

```
Roaming  Essid/Bssid/Phy  Profile  Forward mode  Type  Host Name
---- ----  ---------------  -------  ------------  ----  ---------


User Entries: 1/1
```

Querying a User

This XML request uses the **user_query** command to get the status and details of a user connected to your network.

```
C:\Python27>python ArubaXMLDemo.py --switch-ip=192.0.2.1 --secret=aruba123 --
command=user_query --ip=192.0.2.2
```

The command sends the following information in the XML request to Mobility Master:

- **--switch-ip**: IP address of Mobility Master
- **--secret**: Shared secret key (sent as plain text)
- **--command**: XML API command
- **--ip**: IP address of the user

Mobility Master Response

Mobility Master processes using an XML format and sends the following response to the XML API server.

```
Prepared XML buf
---------------------------------
xml=<aruba command="user_query">
<ipaddr>192.0.2.2</ipaddr>
<authentication>cleartext</authentication>
<key>aruba123</key>
<version>1.0</version>
</aruba>
---------------------------------
Sending XML request to https://192.0.2.1/auth/command.xml
Controller response status: 200
Response XML
-----------------------------------
<aruba>
   <status>Ok</status>
   <code>0</code>
   <macaddr>a4:5e:60:c3:10:59</macaddr>
   <ipaddr>192.0.2.2</ipaddr>
   <name>John</name>
   <role>authenticated</role>
   <type>Wireless</type>
   <vlan>1034</vlan>
   <location>ap225-sales</location>
   <age>00:03:51</age>
   <auth_status>Authenticated</auth_status>
   <auth_server>clearpass-hq1</auth_server>
   <auth_method>802.1X</auth_method>
   <essid>ethersphere-wpa2</essid>
   <bssid>9c:1c:12:92:2e:f1</bssid>
   <phy_type>a-VHT-80</phy_type>
   <mobility_state>Wireless</mobility_state>
   <in_packets>93400</in_packets>
   <in_octets>24947332</in_octets>
   <out_packets>89042</out_packets>
   <out_octets>79397284</out_octets>
</aruba>
```

Mobility Master CLI

The output of the **show user** command displays the client information.

```
(host) #show user

Users
-----
IP           MAC                Name    Role         Age(d:h:m)  Auth           VPN link
----------   ------------       ------  ----         ----------  ----           --------
192.0.2.2    a4:5e:60:c3:10:59  John    authenticate 00:03:51    Authenticated

AP name         Roaming   Essid/Bssid/Phy                              Profile  Forward mode
-------         --------  ---------------                              -------  ------------
ap225-sales     Wireless  ethersphere-wpa2/9c:1c:12:92:2e:f1/a-VHT-80

Type  Host Name
----  ---------

User Entries: 1/1
```

Logging Out a User

This XML request uses the **user_logout** command to revert the user to the initial role. This is only applicable to captive portal users.

```
C:\Python27>python ArubaXMLDemo.py --switch-ip=192.0.2.1 --secret=aruba123 --
command=user_logout --ip=192.0.2.2
```

The command sends the following information in the XML request to Mobility Master:

- **--switch-ip**: IP address of Mobility Master
- **--secret**: Shared secret key (sent as plain text)
- **--command**: XML API command
- **--ip**: IP address of the user

Mobility Master Response

Mobility Master processes using an XML format and sends the following response to the XML API server.

```
Prepared XML buf
--------------------------------
xml=<aruba command="user_logout">
<ipaddr>192.0.2.2</ipaddr>
<authentication>cleartext</authentication>
<key>aruba123</key>
<version>1.0</version>
</aruba>
--------------------------------
Sending XML request to https://192.0.2.1/auth/command.xml
Controller response status: 200
Response XML
-----------------------------------
<aruba>
   <status>Ok</status>
   <code>0</code>
</aruba>
```

Mobility Master CLI

The output of the **show user** command displays the client information.

```
(host) #show user

Users
-----
IP           MAC                Name    Role         Age(d:h:m)  Auth           VPN link
```

```
----------   ------------       ------  ----         ----------  ----            --------
192.0.2.2    a4:5e:60:c3:10:59  John    initial      00:00:06    Unauthenticated

AP name         Roaming   Essid/Bssid/Phy                              Profile  Forward mode
-------         --------  ---------------                              -------  ------------
ap225-sales     Wireless  ethersphere-wpa2/9c:1c:12:92:2e:f1/a-VHT-80

Type  Host Name
----  ---------

User Entries: 1/1
```

Topics in this chapter include:

- Understanding Mode Support on page 1239
- Understanding Basic System Defaults on page 1240
- Understanding Default Management User Roles on page 1248
- Understanding Default Open Ports on page 1251

Understanding Mode Support

Most ArubaOS features are supported in all forwarding modes. However, there are 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.

The following table describes the features that are not supported in each forwarding mode.

Table 275: *Features not Supported in Each Forwarding Mode*

Forwarding Mode	Feature Not Supported
Split Tunnel Mode on Remote APs	■ Bandwidth based CAC ■ DMO ■ DHCP Fingerprinting ■ IGMP Proxy Mobility ■ Layer-2 Mobility ■ Layer-3 Mobility ■ Mobile IP ■ MultiZone ■ Named VLAN ■ TKIP countermeasure management ■ Video over Mesh ■ VLAN pooling ■ Voice over Mesh
Bridge Mode on Campus APs or Remote APs	■ AirGroup ■ AppRF ■ Automatic Voice Flow Classification ■ Bandwidth based CAC ■ Broadcast filter ■ DHCP enforcement

Forwarding Mode	Feature Not Supported
	■ DHCP fingerprint ■ DMO ■ Firewall – Alcatel NOE Support ■ Firewall – SIP, SCCP, RTP, or RTSP Voice Support ■ H.323 ALG ■ IGMP Proxy Mobility ■ Layer 3 Mobility ■ Lync SDN API ■ Management: Voice client statistics ■ Management: Voice client troubleshooting ■ Management: Voice-specific views ■ Mobile IP ■ MPSK ■ Named VLAN ■ NOE ALG ■ Power save: Drop wireless multicast traffic ■ MultiZone ■ Power save: Proxy ARP (global) ■ Power save: Proxy ARP (per-SSID) ■ Power save: Wireless battery boost ■ Rate Limiting for broadcast or multicast ■ SCCP ALG ■ SIP ALG ■ SIP: CAC enforcement enhancements ■ SIP: Delay measurement ■ SIP: Phone number awareness ■ SIP: R-Value computation ■ SIP: SIP authentication tracking ■ Station blacklist ■ Station blacklist by an ACL action ■ SVP ALG ■ TKIP countermeasure management ■ Video over Mesh ■ Vocera ALG ■ Voice over Mesh ■ Voice protocol monitoring or reporting ■ WebCC ■ WPA3 ■ XML-API

Understanding Basic System Defaults

The default administrator user name is **admin**, and the password should be set up during the initial setup dialog. The ArubaOS software includes several predefined network services, firewall policies, and roles.

Predefined Network Services

The following table lists the predefined network services and their protocols and ports.

Table 276: *Predefined Network Services*

Name	Protocol	Port(s)
svc-dhcp	udp	67 68
svc-snmp-trap	udp	162
svc-smb-tcp	tcp	445
svc-https	tcp	443
svc-ike	udp	500
svc-l2tp	udp	1701
svc-syslog	udp	514
svc-pptp	tcp	1723
svc-telnet	tcp	23
svc-sccp	tcp	2000
svc-tftp	udp	69
svc-sip-tcp	tcp	5060
svc-kerberos	udp	88
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

Name	Protocol	Port(s)
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
svc-v6-icmp	icmp	0
any	any	0

Predefined Policies

The following table lists predefined policies.

Table 277: *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 Aruba 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-proxy2 dst-nat 8088`
`user any svc-http-proxy3 dst-nat 8088`
	Enables captive portal authentication.
	1. Any HTTPS traffic destined for the managed device will be routed through NAT to port 8081, where the captive portal server will answer.
	2. All HTTP traffic to any destination will be routed through NAT to the managed device on port 8080, where an HTTP redirect will be issued.
	3. All HTTPS traffic to any destination will be routed through NAT to the managed device on port 8081, where an HTTP redirect will be issued.
	4. All HTTP proxy traffic will be routed through NAT to the managed device 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 managed device on the standard HTTPS port (443) the client will be routed through NAT 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 managed device'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.

Predefined Policy	Description

`ip access-list session ap-acl`
`any any udp 5000 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 managed device.

`ip access-list session validuser` `any any any permit`	This firewall rule controls which users will be added to the user table of the managed device 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 managed device 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 routed through source NAT to the IP address of the managed device.
`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.

Predefined Policy	Description
`ip access-list session dhcp-acl` `any any svc-dhcp permit`	Permits all DHCP traffic. If DHCP is not allowed, clients will not be able to request or renew IP addresses.
`ip access-list session http-acl` `any any svc-http permit`	Permits all HTTP traffic.

`ip access-list session svp-acl`
`any any svc-svp permit queue high`
`user host 224.0.1.116 any permit`
	Use for Spectralink VoIP devices to automatically permit and prioritize Spectralink Voice Protocol.

`ip access-list session noe-acl` `any any svc-noe permit queue high`	Use for Alcatel NOE VoIP devices to automatically permit and prioritize NOE traffic.

`ip access-list session h323-acl`
`any any svc-h323-tcp permit queue high`
`any any svc-h323-udp permit queue high`
	Use for H.323 VoIP devices to automatically permit and prioritize H.323 traffic.

`ipv6 access-list session v6-control`
`user any udp 68 deny`
`any any svc-v6-icmp permit`
`any any svc-v6-dhcp permit`
`any any svc-dns permit`
`any any svc-tftp permit`
	Provides equivalent functionality to the "control" policy, but for IPv6 clients.

`ipv6 access-list session v6-icmp-acl` `any any svc-v6-icmp permit`	Permits all ICMPv6 traffic.
`ipv6 access-list session v6-https-acl` `any any svc-https permit`	Permits all IPv6 HTTPS traffic.
`ipv6 access-list session v6-dhcp-acl` `any any svc-v6-dhcp permit`	Permits all IPv6 DHCP traffic.
`ipv6 access-list session v6-dns-acl` `any any svc-dns permit`	Permits all IPv6 DNS traffic.
`ipv6 access-list session v6-allowall` `any any any permit`	Permits all IPv6 traffic.
`ipv6 access-list session v6-http-acl` `any any svc-http permit`	Permits all IPv6 HTTP traffic.
`ipv6 access-list session v6-tftp-acl` `any any svc-tftp permit`	Permits all IPv6 TFTP traffic.

`ipv6 access-list session v6-logon-control`
`user any udp 68 deny`
`any any svc-v6-icmp permit`
`any any svc-v6-dhcp permit`
`any any svc-dns permit`
	Provides equivalent functionality to the "logon-control" policy, but for IPv6 clients.

Validuser and Logon-control ACLs

Default firewall rules for both the validuser and logon-control ACLs prevent malicious users by blocking self-assigned IPs.

A client with the correct source address can send traffic to the below networks as a destination IP address. The default firewall rules deny traffic FROM the reserved addresses.

The following networks can be blocked by the default firewall rules in both the validuser and logon-control ACLs:

- Network packets where the source address of the network packet is defined as being on a broadcast network (source address == 255.255.255.255)
- Network packets where the source address of the network packet is defined as being on a multicast network (source address = 224.0.0.0 – 239.255.255.255)
- Network packets where the source address of the network packet is defined as being a loopback address (127.0.0.1 through 127.255.255.254)
- Network packets where the source or destination address of the network packet is a link-local address (169.254.0.0/16)
- Network packets where the source or destination address of the network packet is defined as being an address “reserved for future use” as specified in RFC 5735 for IPv4; (240.0.0.0/4)
- Network packets where the source or destination address of the network packet is defined as an “unspecified address”(::/128) or an address “reserved for future definition and use”(addresses other than 2000::/3) as specified in RFC 3513 for IPv6. The IPv6 “an unspecified address”(::/128) is currently being checked in datapath and the packet is dropped. This is the default behavior and you can view the logs by enabling **firewall enable-per-packet-logging** configuration.

Predefined Roles

The following table lists predefined roles.



If you upgrade from a previous ArubaOS release, your existing configuration may have additional or different predefined roles. The information in this section only describes the predefined roles for this release.

Table 278: *Predefined Roles*

Predefined Role	Description

`user-role ap-role`
`session-acl control`
`session-acl ap-acl`
	This is an internal role and should not be edited.

`user-role default-vpn-role`
`session-acl allowall`
`ipv6 session-acl v6-allowall`
	This is the default role used for VPN-connected clients. It is referenced in the default "aaa authentication vpn" profile.

`user-role voice`
`session-acl sip-acl`
`session-acl noe-acl`
`session-acl svp-acl`
`session-acl vocera-acl`
`session-acl skinny-acl`
`session-acl h323-acl`
`session-acl dhcp-acl`
`session-acl tftp-acl`
`session-acl dns-acl`
`session-acl icmp-acl`
	This role can be applied to voice devices in order to automatically permit and prioritize all VoIP protocols.

Predefined Role	Description

`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 or pass through. The logon role should be edited to provide only the required services to a pre-authenticated user. For example, VPN pass through should be disabled if it is not needed.

`user-role <ssid>-logon`
`session-acl control`
`session-acl captiveportal`
`session-acl vpnlogon`
	This role is only generated when creating a new WLAN using the WLAN wizard. The WLAN wizard creates this role when captive portal is enabled and a PEFNG license is installed. This is the initial role that a client will be placed in prior to captive portal authentication. By using a different logon role for each SSID, it is possible to enable multiple captive portal profiles with different customization.

`user-role <ssid>-captiveportal-profile`	When utilizing the WLAN Wizard and you do not have a PEF NG installed and you are configuring an Internal or Guest WLAN with captive portal enabled, the managed device creates an implicit user role with the same name as the captive portal profile, <ssid>-captiveportal-profile.

Predefined Role	Description
	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 managed device, 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.
`sys-ap-role`	This role is for the AP to come with and it has ACLs related to AP termination.
`default-iap-user-role`	This role is for the users coming from IAP VPN tunnel.
`denyall`	This role will deny the user from terminating on the managed device.

Understanding Default Management User Roles

The ArubaOS software includes predefined management user roles.



If you upgrade from a previous ArubaOS release, your existing configuration may have different management roles. The information in this section only describes the predefined management roles for this release.

Table 279: *Predefined Management Roles*

Predefined Role	Permissions
ap-provisioning	This role permits access only to AP provisioning commands and no access to other configuration commands on the Mobility Master.
guest-provisioning	This role permits access to configuring guest users in the managed device'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 managed device using a third-party location appliance, enter: http[s]://<ipaddress>[:port]/screens/wms/wms.login. You are prompted to enter your username and password (for example, the username and password associated with the location API management role). Once authenticated, you can use an API call to request location information from the managed device, for example: http[s]://<ipaddress>[:port]/screens/wms/wms.cgi?opcode=wlm-get-spot&campus-name=<campus id>&building-name<building id>&mac=<client1>,<client2>....

Predefined Role	Permissions
nbapi-mgmt	This role permits configuring a NBAPI management role.
network-operations	**Monitoring** > **Network** > **All Access PointsMonitoring** > **Network** > **All Wired Access Points** You can view the reports created by the following CLI commands: ■ `DB:opcode=monitor-summary` ■ `DB:opcode=cr-load` ■ `DB:opcode=wlm-search&class=probes&start` ■ `DB:opcode=wlm-search&class=amii` ■ `DB:opcode=monitor-get-all-gps&status=any` ■ `show ap-group` ■ `show vlan status` **Monitoring** > **Controller** > **Controller Summary** You can view the reports created by the following CLI commands: ■ `show switches` ■ `show switches summary` **Monitoring** > **Controller** > **Air Monitors** You can view the reports created by the following CLI commands: ■ `show wlan-ap start*` **Monitoring** > **Controller** > **Clients** You can view the reports created by the following CLI commands: ■ `show ip mobile host` ■ `show ip mobile trail {<ipaddr> | <macaddr>}` ■ `<span class="CLI">show ap essid</span>` ■ `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/module/port>` ■ `show interface fastethernet <slot/module/port> counters` ■ `clear counters fastethernet <slot/module/port>` ■ `show snmp trap-queue <page>` **Monitoring** > **Controller** > **Clients** > **Packet CaptureMonitoring** >**Controller** > **Clients** > **LocateMonitoring** > **Controller** > **Clients** > **Debug** You can view the reports created by the following CLI commands: ■ `aaa user debug mac` **Monitoring** > **Controller**> **Clients** > **Disconnect** You can view the reports created by the following CLI commands: ■ `stm kick-off-sta <macaddr>` ■ `aaa user logout <ipaddr>`

Predefined Role	Permissions
network-operations (continued)	**Monitoring** > **Controller**> **Clients** > **Blacklist** You can view the reports created by the following CLI commands: ■ `stm add-blacklist-client <macaddr>` ■ `aaa user delete {<ipaddr> | all | mac <macaddr> | name <username> | role <role>}` **Monitoring** > **Controller** > **Blacklist Clients** You can view the reports created by the following CLI commands: ■ `stm remove-blacklist-client <macaddr>` **Monitoring** > **Controller** > **External Services Interface** You can view the reports created by the following CLI commands: ■ `show esi groups` ■ `show esi servers` ■ `show esi ping` ■ `show esi parser stats` **Monitoring** > **Controller** > **Ports** You can view the reports created by the following CLI commands: ■ `show model-switch-internal`* ■ `show slots` ■ `show private port status`* ■ `show vlan` **Monitoring** > **Controller**> **Inventory** You can view the reports created by executing the following command: ■ `show keys` **Monitoring** > **WLAN** You can view the reports created by the following CLI commands: ■ `DB:opcode=get-permissions` ■ `DB:opcode=cr-load` ■ `show switches` ■ `show switches summary` **Monitoring** > **Voice** You can view the reports created by the following CLI commands: ■ show ap association voip-only ■ `show ap active voip-only` ■ `show voice call-counters` ■ `show voice client status` ■ `show voice call-quality` ■ `show voice call-density` ■ `show voice call-cdrs` ■ `show voice call-perf`
root	This role permits access to all management functions (commands and operations) on the managed device.
read-only	This role permits access to CLI show commands or WebUI monitoring pages only.

Predefined Role	Permissions
standard	This role has root privileges but cannot make changes to the management users. The purpose of creating this role is to prevent changes to the local account from externally authenticated management user.

Understanding Default Open Ports

By default, Aruba managed devices 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 280.

Table 280: *Default (Trusted) Open Ports*

Port Number	Protocol	Where Used	Description
17	TCP	managed device	This is used for certain types of VPN clients that accept a banner (QOTD). During normal operation, this port will only accept a connection and immediately close it.
21	TCP	managed device	
22	TCP	managed device	SSH
23	TCP	AP and managed device	Telnet is disabled by default but the port is still open.
53	UDP	managed device	Internal domain.
67	UDP	AP (and managed device if DHCP server is configured)	DHCP server.
68	UDP	AP (and managed device if DHCP server is configured)	DHCP client.
69	UDP	managed device	TFTP
80	TCP	AP and managed device	Used for remote packet capture where the capture is saved on the access point. Provides access to the WebUI on the managed device.
123	UDP	managed device	NTP
161	UDP	AP and managed device	SNMP. Disabled by default.

Port Number	Protocol	Where Used	Description
443	TCP	managed device	Used internally for captive portal authentication (HTTPS) and is exposed to wireless users. A default self-signed certificate is installed in the managed device. Users in a production environment are urged to install a certificate from a well known CA such as Verisign. Self-signed certs are open to man-in-the-middle attacks and should only be used for testing. **Required for VIA**: During the initializing phase, VIA uses HTTPS connections to perform trusted network and captive portal checks against the managed device. It is mandatory that you enable port 443 on your network to allow VIA to perform these checks.
500	UDP	managed device	ISAKMP
514	UDP	managed device	Syslog
1701	UDP	managed device	L2TP
1723	TCP	managed device	PPTP
2300	TCP	managed device	Internal terminal server opened by **telnet soe** command.
3306	TCP	managed device	Remote wired MAC lookup.

Port Number	Protocol	Where Used	Description
4343, 443	TCP	managed device	HTTPS. Both port 4343 and 443 are supported. If port 4343 is used it redirects to port 443. If port 443 is used it continues to connect using this port. A default self-signed certificate is installed in the managed device. 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	managed device	sae-urn **Required for VIA**: During the initializing phase, VIA uses HTTPS connections to perform trusted network and captive portal checks against the managed device. It is mandatory that you enable port 4500 on your network to allow VIA to perform these checks.
8080	TCP	managed device	Used internally for captive portal authentication (HTTP-proxy). This port is not exposed to wireless users.

Port Number	Protocol	Where Used	Description
8081	TCP	managed device	Used internally for captive portal authentication (HTTPS). Not exposed to wireless users. A default self-signed certificate is installed in the managed device. 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	managed device	Used internally for SSO authentication (HTTP). Exposed to wired and wireless users.
8083	TCP	managed device	Used internally for SSO authentication (HTTPS). Not exposed to wireless users.
8088	TCP	managed device	For internal use.
8200	UDP	managed device	The ADP.
8211	UDP	managed device	For internal use.
8888	TCP	managed device	Used for HTTP access.

This chapter describes how to configure several DHCP vendor-specific options.

Topics in this chapter include:

- Configuring a Windows-Based DHCP Server on page 1255
- Enabling DHCP Relay Agent Information Option (Option 82) on page 1257
- Enabling Linux DHCP Servers on page 1261

Configuring a Windows-Based DHCP Server

Configuring a Microsoft Windows-based DHCP server to send Option 43 to the DHCP client on an Aruba AP consists of the following two tasks:

- Configuring DHCP Response (Option 43)
- Configuring DHCP Vendor Class Request (Option 60)

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 Information (VSI), also called Option 43. When a client or an AP requests for Option 43 (VSI), the DHCP server responds with the IP address of the managed device configured by the administrator in the DHCP pool.

When a factory-default AP boots up and requests an IP address, the AP includes a Vendor Class Identifier (VCI) string by default in its DHCP request, also called Option 60. VCI is a text string that uniquely identifies a type of vendor device. Based on the VCI string, the DHCP server responds with the correct VSI included in Option 43.

Configuring DHCP Response (Option 43)

Configuring DHCP response Option 43 returns the IP address of the Aruba managed device to an Aruba DHCP client. This information allows Aruba APs to auto-discover the Mobility Master and obtain their configuration.

Configuring Option 43 Using the Windows DHCP Server

The following procedure configures DHCP response (Option 43) using the Windows DHCP server:

1. On the DHCP server, navigate to **Start > Administration Tools > DHCP** to open the DHCP server administration tool.
2. Find your server and right-click on the scope to be configured under the server name.
3. Click on the **Scope Options** entry and select **Configure Options**.
4. In the **Scope Options** window, scroll down and select **043 Vendor Specific Info**.
5. In the **Data Entry** field, click anywhere in the area under the ASCII heading and enter ASCII : Loopback address of the managed device.

6. Click **OK** to save the configuration.

   Option 43 is configured for this DHCP scope.



Though you entered the IP address of the managed device in ASCII text, the IP address is displayed in binary form.

Configuring DHCP Vendor Class Request (Option 60)

When an AP sends a DHCP request, the AP identifies itself to the DHCP server by setting its VCI to **ArubaAP** in the DHCP request. Configuring DHCP vendor class request (Option 60) consists of the following two tasks:

1. Creating a new vendor class—Configure a Vendor Class Identifier (VCI) on the DHCP server for each AP. This VCI must match the VCI (**ArubaAP**) defined in Option 60 on the AP.
2. Creating a new policy—Create a new policy and assign server values in Option 43 for the newly-created vendor class.



You must configure the DHCP vendor class only when the other devices in the same DHCP scope use DHCP Option 43, or when you want to display the IP address of the managed device to APs only.

Creating a New Vendor Class

The following procedure configures a new vendor class on the Windows-based DHCP server:

1. On the DHCP server, navigate to **Start** > **Administration Tools** > **DHCP** to open the DHCP server administration tool.
2. Find your server and right-click on **IPv4**.
3. Click **Define Vendor Classes**.

   The **New Class** window is displayed.
4. In the **New Class** window, enter the following information:
   - **Display name**—ArubaAP
   - **Description**—VCI: ArubaAP
5. In the **ASCII** column, enter **ArubaAP** as the VCI.
6. Click **OK**.

   The new vendor class (**ArubaAP**) is displayed in the **DHCP Vendor Classes** window.

Creating a New Policy

The following procedure configures a new policy on the Windows-based DHCP server:

1. On the DHCP server, navigate to **Start** > **Administration Tools** > **DHCP** to open the DHCP server administration tool.
2. Find your server and right-click on **Policies**.
3. Click **New Policy**.

   The **DHCP Policy Configuration Wizard** window is displayed.

4. Under **General** tab, enter the following information:
   - **Policy Name**—ArubaAP
   - **Description**—Aruba Access Point to controller
5. Under **Conditions** tab, perform the following steps:
   a. Select **Vendor Class** from the **Criteria** drop-down list.
   b. Select **Equals** from the **Operator** drop-down list.
   c. Select the newly-created vendor class (**ArubaAP**) from the **Value** drop-down list, and click **Add**.
6. Under **Options** tab, perform the following steps:
   a. Select **DHCP Standard Options** from the **Vendor class** drop-down list.
   b. Select **043 Vendor Specific Info** check box from the **Available Options** column.
   c. In the **Data Entry** field, click anywhere in the area under the **ASCII** column and enter the IP address of your managed device that the AP should connect to.
   d. Click **Apply**, and then click **OK**.

   The new policy is applied and the DHCP server values are assigned in Option 43 for the newly-created vendor class (**ArubaAP**).

   ### Important Points to Note

   - Repeat the above procedure to create more policies for various device types that require Option 43 or different options.
   - If you have multiple Windows-based DHCP servers IP addresses for your APs, you must configure DHCP vendor class on each of them.
   - You can create new policies on a per scope level, or for all the IPv4 scopes based on your deployment.

Configuring Option 52 Using the Windows DHCP Server

Starting from ArubaOS 8.1.0.0, APs can discover a managed device in an IPv6 deployment using DHCPv6 Option 52. The AP as a DHCPv6 client requests for Option 52 in a parameter request list of Solicit and Request DHCPv6 packets. If the DHCPv6 server has configured this option, the DHCPv6 server will return this option to the AP in the Advertise and Reply packet, the AP will then parse managed device address from option 52 and try to connect to it.

The following procedure configures Option 52 using the Windows DHCP server:

1. On the DHCP server, navigate to **Start > Administration Tools > DHCP** to open the DHCP server administration tool.
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, scroll down and select **052 capwap-ac-v6**.
4. In the **Data Entry** field, click anywhere in the area under the ASCII heading and enter ASCII : Master Controller IP address of the master controller.
5. Click **OK** to save the configuration.

Option 52 is configured for this DHCP scope.

Enabling DHCP Relay Agent Information Option (Option 82)

Option-82 feature allows the DHCP Relay Agent to insert specific information into a client request that is being forwarded to a DHCP server. Option-82 can be customized to cater to the requirements of any ISP to make access control decisions using the Arubamanaged device.

The managed device, when acting as a DHCP relay agent can be configured with the following sub-type options,

- Type 1 – Circuit ID (AP, Port)
- Type 2 – Remote ID (Client SSID)
- Type 5 – Link Selection (Local VLAN Network)

Starting from ArubaOS 8.1.0.0, a XML definition file has been introduced to provide flexibility to configure multiple sub-type options. The XML file is used as the input from the user and is validated against an XSD file stored under flash on the managed device. The format in the XML file is parsed and stored by the DHCP relay agent module to ensure option-82 related values are inserted in the DHCP request packets from the client to the server.

Starting from ArubaOS 8.1.0.0, when IP Helper is enabled on L3 interface, DHCP discover broadcast is filtered at the datapath level and is unicast to the configured helper device. DHCP Option-82 on L2 VLAN can now be enabled without the helper address.



DHCP Option-82 is supported only for wired and wireless IPv4 clients and applicable to wireless clients terminating in AP Tunnel and D-Tunnel modes.

Sample XML Format

The following is a sample XML file which specifies DHCP Option-82 circuit and remote IDs and link selection fields.

```
<?xml version="1.0" encoding="UTF-8"?>
   <dhcpopt82>
     <circuit_id>
       <param>
         <type>var</type>
           <val>apmac</val>
           <delim>-</delim>
       </param>
      </circuit_id>
      <remote_id>
       <param>
        <type>var</type>
          <val>cmac</val>
          <delim>:</delim>
        </param>
      </remote_id>
   <link_selection>
    <param>
     <type>var</type>
      <val>vlanip</val>
     </param>
   </link_selection>
</dhcpopt82>
```

The following table describes the fields to be configured for wired and wireless clients.

Type of client	Fields to be configured
Wired clients	link_selection field only
Wireless clients	Any one of the following combinations should be configured if DHCP option 82 is enabled on a VLAN, • circuit_id, remote_id and Link _selection • circuit_id and remote_id • link_selection

The table below lists the elements introduced in the **param** sub-options of the **Circuit ID** and **Remote ID and Link selection** fields:

XML File Parameters

Parameter	Description
Type	Listed below are the types available: ■ var - A DHCP option-82 allowed keyword ■ hex - A hexadecimal string with a maximum of 60 characters ■ str - An ASCII string that with a maximum of 60 characters
Val	This field contains either a hexadecimal string or ASCII string limited to 60 characters, if the type is hex or str. If the type is var then one of the following DHCP option-82 keywords: ■ apname - AP name ■ apmac - AP MAC ■ cmac - Client MAC ■ essid - ESSID broadcasted by the AP ■ bssid - BSSID of the AP ■ vlanip- Local vlan interface
Delim	The delim option is available only for mac based keywords - apmac, cmac, and bssid. The delim field is used if mac addresses are required to be in ASCII format with octets separated with specified ASCII character in the val part of delim field. By default the ASCII MAC separated by a delimiter, will be in lower case. If the user wants to use upper case then the respective mac based val keywords should be written in upper case in XML file. For example: <param> <type>var</type> <val>CMAC</val> <delim>:</delim> </param>

Configuring XML Based DHCP Option-82 Specification

The following procedure is the sequence of steps to be followed if a user wants to use XML based DHCP option-82:

1. From the Mobility Master upload the XML file containing Option-82 specification to flash using the copy command. For example:

```
(host) #copy scp: 10.20.22.42 piddalagi /home/piddalagi/my_dhcp_option_82.xml flash:
my_dhcp_option_82.xml
Password:**********
Secure File Copy:....
```

2. In the configuration terminal execute the following command:

```
#ip dhcp option-82 <xml-file-name-in-flash>
```

For example:

```
(host) ^[md] (config) #ip dhcp option-82 my_dhcp_option_82.xml
```



If you are re-applying a modified XML file with the same file name, ensure that you execute the **no ip dhcp option-82** command before executing the **ip dhcp option-82 <xml-file>** command , for the changes to take effect.

1. After the **ip dhcp option-82 <xml-file-name-in-flash>** command is executed successfully, in the WLAN interface configuration execute **option-82** command to apply option-82 configuration to the DHCP packets that need to be relayed from that interface.

For example:

```
(host) ^[md] (config) #interface vlan 25
(host) ^[md] (config-subif)#option-82
```

2. Execute the following command if the user wants to remove option-82 configuration:

```
(host) (config) #no ip dhcp option-82
```



This command is not successful if there is a WLAN VLAN interface configured with **option-82** command. To remove option-82 configuration, go to the respective VLAN interface and execute **no option-82** command and then execute the **no ip dhcp option-82** command.

Configuring Option 82

The following procedure configures Option 82:

1. In the **Managed Network** node device, navigate to the **Configuration** > **Interfaces** > **VLANs** tab.
2. Select the VLAN ID for which you want to configure Option 82. The **VLANs** table is displayed.
3. From the **VLANs >** table select the VLAN ID again.
4. Select **More** tab from the table that is displayed.
5. Expand the **Other Options** accordion.
6. Select the check box **DHCP Server Option 82** to configure Option 82.
7. Click **Submit**.
8. Click **Pending Changes**.
9. In the **Pending Changes** window, select the check box and click **Deploy changes.**

In the CLI

The following commands configure DHCP Option-82, from device node:

```
(host) [mynode] (config) #show configuration devices
Provisioned Devices
-------------------
Device             Model  Nodepath  VPN Concentrator  Config status
------             -----  --------  ----------------  -------------
00:0c:29:88:5c:17  MC-VA  /md/blr   None              Loaded
Total Devices: 1
```

Execute the following command to enable L3 option-82:

```
(host) [mynode] (config) #cd /md/blr/00:0c:29:88:5c:17
(host) [00:0c:29:88:5c:17] (config) #interface vlan 1
```

```
(host) [00:0c:29:88:5c:17] (config-submode)#op
operstate               Interface Operation state
option-82               Turn on option 82
(SP-MM-110) [00:0c:29:88:5c:17] (config-submode)#option-82
```

Execute the following command to enable L2 option-82

```
(host) [00:0c:29:88:5c:17] (config-submode)#exit
(host) [00:0c:29:88:5c:17] (config) #vlan 1
option-82               Turn on Option-82
```

Execute the following command to remove Option-82 configuration:

```
(host) [mynode] (config) #no ip dhcp option-82
```

Execute the following command to remove DHCP option-82 configuration on L3 VLAN:

```
(host) [mynode] (config) #interface vlan
(host) [00:0c:29:88:5c:17] (config) #no vlan option-82
```

Enabling Linux DHCP Servers

The following is an example configuration for the Linux dhcpd.conf file. After you enter the configuration, you must restart the DHCP service.

```
option serverip code 43 = ip-address;
option serverip code 43 = ip-address;
class "vendor-class"
{
match option vendor-class-identifier;
}
subclass "vendor-class" "ArubaAP"
{
option vendor-class-identifier "ArubaAP";
}
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";
#
#option serverip <loopback-IP-address-of-master-controller>
#
option serverip 10.200.10.10;
range 10.200.10.200 10.200.10.252;
}
```

Chapter 51
802.1X Configuration for IAS and Windows Clients

This chapter provides examples of how to configure a Microsoft Internet Authentication Server, and a Windows XP wireless client for 802.1X authentication with the Mobility Master. For more information on 802.1X Authentication, see 802.1X Authentication on page 268.

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.mspx#EQGAC.

This chapter describes the following topics:

- Configuring Microsoft IAS on page 1262
- Configuring Management Authentication Using IAS on page 1264
- Window XP Wireless Client Sample Configuration on page 1267

Configuring Microsoft IAS

Microsoft IAS provides authentication functions for wireless networks. IAS implements the RADIUS protocol, which is used between the Aruba Mobility Master 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 Aruba Mobility Master as a RADIUS client.



The steps to perform this task may very depending on the version of Windows currently running on your server. For complete details on configuring Windows IAS, refer to the Windows documentation available at www.microsoft.com/downloads.

The following procedure configures a RADIUS client:

1. From your Windows server, navigate to **Start > Settings > Control Panel > Administrative Tools > Internet Authentication Service**.
2. In the **Internet Authentication Service** window, select **RADIUS Clients**.
3. To configure a RADIUS client, select **Action > New RADIUS Client** from the menu at the top of the window.
4. In the **New RADIUS Client** dialog window, enter the name and IP address for the Mobility Master. Click **Next**.
5. Enter and confirm a shared secret.
6. The shared secret is configured on both the RADIUS server and client, and ensures that an unauthorized client cannot perform authentication against the server.
7. 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 the Active Directory. These policies are responsible for passing group information back to the Mobility Master 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 Mobility Master configuration shown in Example Configurations on page 281:

- 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 or password is supplied, the user’s or computer’s remote access permission is set to “Allow”.

The following procedure configures a policy:

1. In the **Internet Authentication Service** window, select **Remote Access Policies**.
2. Navigate to **Action > New Remote Access Policy** to add a new policy.
3. Click **Next** on the initial wizard window to proceed.
4. Enter a name for the policy, for example, Wireless Computers and click **Next**.
5. In the **Access Method** window, select the **Wireless** option, then click **Next**.
6. In the **User or Group Access** window, select **Group** and click **Add** to add the group of users to which this policy applies (for example, “Domain Computers”). Click **Next**.

7. For **Authentication Methods**, select either **Protected EAP (PEAP)** or **Smart Card** or **Other Certificate**.
8. Click **Configure** to select additional properties.
9. Select a server certificate.
10. 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 CA and installed on the IAS system. On each wireless client device, the local CA is added as a trusted CA, thus allowing this certificate to be trusted.
11. For PEAP, select the inner authentication method.
12. The authentication method shown is MS-CHAPv2. This should be the only EAP authentication type that should be selected as password authentication is being used on this network.
13. 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.
14. Click **OK**.

Configuring RADIUS Attributes

In the configuration example for 802.1X, the Mobility Master restricts network access privileges based on the group membership of the computer or user. In order for this to work, the Mobility Master must be told to which group the user belongs. This is accomplished using RADIUS attributes returned by the authentication server.

The following procedures configure RADIUS attributes:

1. In the **Internet Authentication Service** window, select **Remote Access Policies**.
2. Open the remote access policy you want to configure, and select the **Advanced** tab.
3. Click **Add** to configure an attribute.
4. Select the **Class** attribute.
5. Enter the value for this attribute. For example, for the **Wireless-Computers** policy, the **Class** attribute returned to the Mobility Master should contain the value “computer”.
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” after successful completion.

Configuring Management Authentication Using IAS

Before you can configure the Mobility Master for management authentication using Windows IAS, you must perform the following steps to configure a Windows IAS RADIUS server on your Windows client.



The steps to perform this task may very depending on the version of Windows currently running on your server. For complete details on configuring Windows IAS, refer to the Windows documentation available at 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 is displayed.

2. Verify if IAS is running.
3. If IAS is running, a green arrow icon is displayed 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.
4. From the **Internet Authentication Service** window, right click the **RADIUS Clients** folder and select **New RADIUS Client**. The **New RADIUS Client** window is displayed.
5. Enter a name for the RADIUS client and enter the Mobility Master's IP address or DNS name. Click **Next**.
6. Enter and confirm the shared secret key for the Mobility Master. Click **Finish**.

Creating a Remote Policy

The following procedure describes the steps to create a remote policy:

1. From the **Internet Authentication Service** window, right click the **Remote Access Policies** folder and select **New Remote Access Policy**.
2. The **New Remote Access Policy** wizard is displayed. 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. 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**.

Defining Properties for Remote Policy

The following procedure describes the steps to define properties for remote policy:

1. In the **Internet Authentication Service** window, click the **Remote Access Policy** icon. All configured remote access policies are displayed in the right window pane.
2. Right-click the policy you just created, and select **Properties**. The **Properties** window is displayed.
3. Select the **Grant remote access permission** option, and click **Edit Profile**. The **Edit Profile** window is displayed.
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 is displayed.
8. Scroll down the list of attributes and select **Vendor-Specific**, then click **Add**. The **MultiValued Attribute Information** window is displayed.
9. Click **Add**.
10. Enter the vendor code **14823** and select the option **Yes, it conforms**.
11. Click **Configure Attribute**. The **Configure VSA** window is displayed.
12. In the **Vendor-assigned attribute number** field, enter **3**.
13. In the **Attribute value** field, enter **7**.

14. Click **OK** to save the settings.
15. Click **Apply**.
16. Click **Apply**.

Now that you have defined your remote policy properties, you must create a user entry in the Windows active directory. The steps to complete this process will vary, depending on the version of Windows currently running on your server. The procedure below should be used only as a guideline.

Creating a User Entry in Windows Active Directory

The following procedure describes the steps to create a user entry in Windows Active Directory:

1. Open the **Active Directory Users and Computers** tool on your Windows server.
2. Create a new user entry on the Windows Active directory.
3. Once you have created the new user, right-click the user name and select **Properties**.
4. Click the **Dial-in** tab and select **"Allow access"** for the user.
5. Click **OK** to save your settings.

Configure the Mobility Master to Use IAS Management Authentication

The following procedure describes the steps to configure the Mobility Master to user IAS management authentication.

1. In the **Mobility Master** node hierarchy, navigate to **Configuration > Authentication > Auth Servers**.
2. Click **+** in **All Servers** table. The **New Server** dialog box is displayed.
3. Enter a value for following:
   a. Name
   b. IP address
   c. From the **Type** drop-down list, select **Radius**.
4. Click **Submit**.
5. Select the RADIUS server that you created.
6. Enter and then retype the shared key for the server.
7. Click **Submit**.
8. Click **+** in **Server Groups** table. The **Add Server Group** dialog box is displayed.
9. Enter a name for the server group.
10. Click **Submit**.
11. Select the server group that you created.
12. Click **+** in **Server** table. A dialog box is displayed that prompts you to select a RADIUS server.
13. Ensure that the **Add existing server** option is selected.
14. Select the RADIUS server you created. Click **Submit.**
15. Click **Pending Changes**.
16. In the **Pending Changes** window, select the check box and click **Deploy changes**.

Verify Communication Between the Mobility Master and RADIUS Server

After you have configured your Windows Server and the Mobility Master for Windows IAS Management Authentication, you can verify that the Mobility Master and server are communicating using the following procedure:

1. In the **Mobility Master** node hierarchy, navigate to **Diagnostics > Tools > AAA Server Test**.
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 **Test**.

If the Mobility Master displays **Authentication Successful**, then the Mobility Master is able to communicate with the RADIUS server.

Window XP Wireless Client Sample Configuration

This section shows an example of how to configure a Windows XP wireless client using Windows XP's Wireless Zero Configuration service.



The following steps apply to a computer running Windows XP Professional Version 2002 with Service Pack 2. To configure a wireless client on other Windows platforms, see your Microsoft Windows documentation.

1. On the desktop, right-click **My Network Places** and select **Properties**.
2. In the **Network Connections** window, right-click on **Wireless Network Connection** and select **Properties**.
3. Select the **Wireless Networks** tab. This screen displays the available wireless networks and the list of preferred networks. Windows connects to the preferred networks in the order in which they appear in the list.
4. Click the **Advanced** button to display the **Networks to access** window.
5. 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**.
6. In the **Wireless Networks** tab, click **Add** to add a wireless network.
7. Click the **Association** tab to enter the network properties for the SSID.



- This tab configures the authentication and encryption used between the wireless client and the Aruba user-centric network. Therefore, the settings for the SSID that you configure on the client must match the configuration for the SSID on the Mobility Master.

- 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 PSK
- 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 PSK



Do not select the option “This is a computer-to-computer (ad hoc) network; wireless access points are not used”.

8. 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.
   - Select Enable IEEE 802.1X authentication for this network.
   - Select PEAP for the EAP type.
   - Select Authenticate as computer when computer information is available. The client will perform computer authentication when a user is not logged in.
   - Do not select Authenticate as guest when user or computer information is unavailable. The client will not attempt to authenticate as a guest.
   - Select Validate server certificate. This instructs the client to check the validity of the server certificate from an expiration, identity, and trust perspective.
   - Select the trusted Certification Authority 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.
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.

Chapter A
Glossary of Terms

The following table provides a brief description of the terminology used in this guide.

3DES

Triple Data Encryption Standard. 3DES is a symmetric-key block cipher that applies the DES cipher algorithm three times to each data block.

3G

Third Generation of Wireless Mobile Telecommunications Technology. See W-CDMA.

3GPP

Third Generation Partnership Project. 3GPP is a collaborative project aimed at developing globally acceptable specifications for third generation mobile systems.

4G

Fourth Generation of Wireless Mobile Telecommunications Technology. See LTE.

802.11

802.11 is an evolving family of specifications for wireless LANs developed by a working group of the Institute of Electrical and Electronics Engineers (IEEE). 802.11 standards use the Ethernet protocol and Carrier Sense Multiple Access with collision avoidance (CSMA/CA) for path sharing.

802.11 bSec

802.11 bSec is an alternative to 802.11i. The difference between bSec and standard 802.11i is that bSec implements Suite B algorithms wherever possible. Notably, Advanced Encryption Standard-Counter with CBC-MAC is replaced by Advanced Encryption Standard - Galois/Counter Mode, and the Key Derivation Function (KDF) of 802.11i is upgraded to support SHA-256 and SHA-384.

802.11a

802.11a provides specifications for wireless systems. Networks using 802.11a operate at radio frequencies in the 5 GHz band. The specification uses a modulation scheme known as orthogonal frequency-division multiplexing (OFDM) that is especially well suited to use in office settings. The maximum data transfer rate is 54 Mbps.

802.11ac

802.11ac is a wireless networking standard in the 802.11 family that provides high-throughput WLANs on the 5 GHz band.

802.11b

802.11b is a WLAN standard often called Wi-Fi and is backward compatible with 802.11. Instead of the Phase-Shift Keying (PSK) modulation method used in 802.11 standards, 802.11b uses Complementary Code Keying (CCK) that allows higher data speeds and makes it less susceptible to multipath-propagation interference. 802.11b operates in the 2.4 GHz band and the maximum data transfer rate is 11 Mbps.

802.11d

802.11d is a wireless network communications specification for use in countries where systems using other standards in the 802.11 family are not allowed to operate. Configuration can be fine-tuned at the Media Access Control (MAC) layer level to comply with the rules of the country or district in which the network is to be used. Rules are subject to variation and include allowed frequencies, allowed power levels, and allowed signal bandwidth. 802.11d facilitates global roaming.

802.11e

802.11e is an enhancement to the 802.11a and 802.11b specifications that enhances the 802.11 Media Access Control layer with a coordinated Time Division Multiple Access (TDMA) construct. It adds error-correcting mechanisms for delay-sensitive applications such as voice and video. The 802.11e specification provides seamless interoperability between business, home, and public environments such as airports and hotels, and offers all subscribers high-speed Internet access with full-motion video, high-fidelity audio, and VoIP.

802.11g

802.11g offers transmission over relatively short distances at up to 54 Mbps, compared with the 11 Mbps theoretical maximum of 802.11b standard. 802.11g employs Orthogonal Frequency Division Multiplexing (OFDM), the modulation scheme used in 802.11a, to obtain higher data speed. Computers or terminals set up for 802.11g can fall back to speed of 11 Mbps, so that 802.11b and 802.11g devices can be compatible within a single network.

802.11h

802.11h is intended to resolve interference issues introduced by the use of 802.11a in some locations, particularly with military Radar systems and medical devices. Dynamic Frequency Selection (DFS) detects the presence of other devices on a channel and automatically switches the network to another channel if and when such signals are detected. Transmit Power Control (TPC) reduces the radio frequency (RF) output power of each network transmitter to a level that minimizes the risk of interference.

802.11i

802.11i provides improved encryption for networks that use 802.11a, 802.11b, and 802.11g standards. It requires new encryption key protocols, known as Temporal Key Integrity Protocol (TKIP) and Advanced Encryption Standard (AES).

802.11j

802.11j is a proposed addition to the 802.11 family of standards that incorporates Japanese regulatory extensions to 802.11a; the main intent is to add channels in the radio frequency (RF) band of 4.9 GHz to 5.0 GHz.

802.11k

802.11k is an IEEE standard that enables APs and client devices to discover the best available radio resources for seamless BSS transition in a WLAN.

802.11m

802.11m is an Initiative to perform editorial maintenance, corrections, improvements, clarifications, and interpretations relevant to documentation for 802.11 family specifications.

802.11n

802.11n is a wireless networking standard to improve network throughput over the two previous standards, 802.11a and 802.11g. With 802.11n, there will be a significant increase in the maximum raw data rate from 54 Mbps to 600 Mbps with the use of four spatial streams at a channel width of 40 MHz.

802.11r

802.11r is an IEEE standard for enabling seamless BSS transitions in a WLAN. 802.11r standard is also referred to as Fast BSS transition.

802.11u

802.11u is an amendment to the IEEE 802.11 WLAN standards for connection to external networks using common wireless devices such as smartphones and tablet PCs. The 802.11u protocol provides wireless clients with a streamlined mechanism to discover and authenticate to suitable networks, and allows mobile users to roam between partner networks without additional authentication. An 802.11u-capable device supports the Passpoint technology from the Wi-Fi Alliance Hotspot 2.0 R2 Specification that simplifies and automates access to public Wi-Fi.

802.11v

802.11v is an IEEE standard that allows client devices to exchange information about the network topology and RF environment. This information is used for assigning best available radio resources for the client devices to provide seamless connectivity.

802.1Q

802.1Q is an IEEE standard that enables the use of VLANs on an Ethernet network. 802.1Q supports VLAN tagging.

802.1X

802.1X is an IEEE standard for port-based network access control designed to enhance 802.11 WLAN security. 802.1X provides an authentication framework that allows a user to be authenticated by a central authority.

802.3af

802.3af is an IEEE standard for Power over Ethernet (PoE) version that supplies up to 15.4W of DC power. See PoE.

802.3at

802.3at is an IEEE standard for PoE version that supplies up to 25.5W of DC power. See PoE+.

A-MPDU

Aggregate MAC Protocol Data Unit. A-MPDU is a method of frame aggregation, where several MPDUs are combined into a single frame for transmission.

A-MSDU

Aggregate MAC Service Data Unit. A-MSDU is a structure containing multiple MSDUs, transported within a single (unfragmented) data MAC MPDU.

AAA

Authentication, Authorization, and Accounting. AAA is a security framework to authenticate users, authorize the type of access based on user credentials, and record authentication events and information about the network access and network resource consumption.

ABR

Area Border Router. ABR is used for establishing connection between the backbone networks and the Open Shortest Path First (OSPF) areas. ABR is located near the border of one or more OSPF areas.

AC

Access Category. As per the IEEE 802.11e standards, AC refers to various levels of traffic prioritization in Enhanced Distributed Channel Access (EDCA) operation mode. The WLAN applications prioritize traffic based on the Background, Best Effort, Video, and Voice access categories. AC can also refer to Alternating Current, a form of electric energy that flows when the appliances are plugged to a wall socket.

ACC

Advanced Cellular Coexistence. The ACC feature in APs enable WLANs to perform at peak efficiency by minimizing interference from 3G/4G/LTE networks, distributed antenna systems, and commercial small cell/femtocell equipment.

Access-Accept

Response from the RADIUS server indicating successful authentication and containing authorization information.

Access-Reject

Response from RADIUS server indicating that a user is not authorized.

Access-Request

RADIUS packet sent to a RADIUS server requesting authorization.

Accounting-Request

RADIUS packet type sent to a RADIUS server containing accounting summary information.

Accounting-Response

RADIUS packet sent by the RADIUS server to acknowledge receipt of an Accounting-Request.

ACE

Access Control Entry. ACE is an element in an ACL that includes access control information.

ACI

Adjacent Channel Interference. ACI refers to interference or interruptions detected on a broadcasting channel, caused by too much power on an adjacent channel in the spectrum.

ACL

Access Control List. ACL is a common way of restricting certain types of traffic on a physical port.

Active Directory

Microsoft Active Directory. The directory server that stores information about a variety of things, such as organizations, sites, systems, users, shares, and other network objects or components. It also provides authentication and authorization mechanisms, and a framework within which related services can be deployed.

ActiveSync

Mobile data synchronization app developed by Microsoft that allows a mobile device to be synchronized with either a desktop or a server running compatible software products.

ad hoc network

An ad hoc network is a network composed of individual devices communicating with each other directly. Many ad hoc networks are Local Area Networks (LANs) where computers or other devices are enabled to send data directly to one another rather than going through a centralized access point.

ADO

Active X Data Objects is a part of Microsoft Data Access Components (MDACs) that enables client applications to access data sources through an (Object Linking and Embedding Database) OLE DB provider. ADO supports key features for building client-server and Web-based applications.

ADP

Aruba Discovery Protocol. ADP is an Aruba proprietary Layer 2 protocol. It is used by the APs to obtain the IP address of the TFTP server from which it downloads the AP boot image.

AES

Advanced Encryption Standard. AES is an encryption standard used for encrypting and protecting electronic data. The AES encrypts and decrypts data in blocks of 128 bits (16 bytes), and can use keys of 128 bits, 192 bits, and 256 bits.

AIFSN

Arbitrary Inter-frame Space Number. AIFSN is set by the AP in beacon frames and probe responses. AIFS is a method of prioritizing a particular category of traffic over the other, for example prioritizing voice or video messages over email.

AirGroup

The application that allows the end users to register their personal mobile devices on a local network and define a group of friends or associates who are allowed to share them. AirGroup is primarily designed for colleges and other institutions. AirGroup uses zero configuration networking to allow Apple mobile devices, such as the AirPrint wireless printer service and the AirPlay mirroring service, to communicate over a complex access network topology.

AirWave Management Client

AirWave Management Client is a Windows software utility that enables client devices (such as a laptop) to act as passive RF sensors and augments the AirWave RAPIDS module.

ALE

Analytics and Location Engine. ALE gives visibility into everything the wireless network knows. This enables customers and partners to gain a wealth of information about the people on their premises. This can be very important for many different verticals and use cases. ALE includes a location engine that calculates associated and unassociated device location periodically using context streams, including RSSI readings, from WLAN controllers or Instant clusters.

ALG

Application Layer Gateway. ALG is a security component that manages application layer protocols such as SIP, FTP and so on.

AM

Air Monitor. AM is a mode of operation supported on wireless APs. When an AP operates in the Air Monitor mode, it enhances the wireless networks by collecting statistics, monitoring traffic, detecting intrusions, enforcing security policies, balancing wireless traffic load, self-healing coverage gaps, and more. However, clients cannot connect to APs operating in the AM mode.

AMON

Advanced Monitoring. AMON is used in Aruba WLAN deployments for improved network management, monitoring and diagnostic capabilities.

AMP

AirWave Management Platform. AMP is a network management system for configuring, monitoring, and upgrading wired and wireless devices on your network.

ANQP

Access Network Query Protocol. ANQP is a query and a response protocol for Wi-Fi hotspot services. ANQP includes information Elements (IEs) that can be sent from the AP to the client to identify the AP network and service provider. The IEs typically include information about the domain name of the AP operator, the IP addresses available at the AP, and information about potential roaming partners accessible through the AP. If the client responds with a request for a specific IE, the AP will send a Generic Advertisement Service (GAS) response frame with the configured ANQP IE information.

ANSI

American National Standards Institute. It refers to the ANSI compliance standards for products, systems, services, and processes.

API

Application Programming Interface. Refers to a set of functions, procedures, protocols, and tools that enable users to build application software.

app

Short form for application. It generally refers to the application that is downloaded and used on mobile devices.

ARM

Adaptive Radio Management. ARM dynamically monitors and adjusts the network to ensure that all users are allowed ready access. It enables full utilization of the available spectrum to support maximum number of users by intelligently choosing the best RF channel and transmit power for APs in their current RF environment.

ARP

Address Resolution Protocol. ARP is used for mapping IP network address to the hardware MAC address of a device.

Aruba Activate

Aruba Activate is a cloud-based service that helps provision your Aruba devices and maintain your inventory. Activate automates the provisioning process, allowing a single IT technician to easily and rapidly deploy devices throughout a distributed enterprise network.

ASCII

American Standard Code for Information Interchange. An ASCII code is a numerical representation of a character or an action.

B-RAS

Broadband Remote Access Server. A B-RAS is a server that facilitates and converges traffic from multiple Internet traffic resources such as cable, DSL, Ethernet, or Broadband wireless.

band

Band refers to a specified range of frequencies of electromagnetic radiation.

BGP

Border Gateway Protocol. BGP is a routing protocol for exchanging data and information between different host gateways or autonomous systems on the Internet.

BLE

Bluetooth Low Energy. The BLE functionality is offered by Bluetooth® to enable devices to run for long durations with low power consumption.

BMC

Beacon Management Console. BMC manages and monitors beacons from the BLE devices. The BLE devices are used for location tracking and proximity detection.

BPDU

Bridge Protocol Data Unit. A BPDU is a data message transmitted across a local area network to detect loops in network topologies.

BRE

Basic Regular Expression. The BRE syntax standards designed by the IEEE provides extension to the traditional Simple Regular Expressions syntax and allows consistency between utility programs such as grep, sed, and awk.

BSS

Basic Service Set. A BSS is a set of interconnected stations that can communicate with each other. BSS can be an independent BSS or infrastructure BSS. An independent BSS is an ad hoc network that does not include APs, whereas the infrastructure BSS consists of an AP and all its associated clients.

BSSID

Basic Service Set Identifier. The BSSID identifies a particular BSS within an area. In infrastructure BSS networks, the BSSID is the MAC address of the AP. In independent BSS or ad hoc networks, the BSSID is generated randomly.

BYOD

Bring Your Own Device. BYOD refers to the use of personal mobile devices within an enterprise network infrastructure.

CA

Certificate Authority or Certification Authority. Entity in a public key infrastructure system that issues certificates to clients. A certificate signing request received by the CA is converted into a certificate when the CA adds a signature generated with a private key. See digital certificate.

CAC

Call Admission Control. CAC regulates traffic volume in voice communications. CAC can also be used to ensure or maintain a certain level of audio quality in voice communications networks.

CALEA

Communications Assistance for Law Enforcement Act. To comply with the CALEA specifications and to allow lawful interception of Internet traffic by the law enforcement and intelligence agencies, the telecommunications carriers and manufacturers of telecommunications equipment are required to modify and design their equipment, facilities, and services to ensure that they have built-in surveillance capabilities.

Campus AP

Campus APs are used in private networks where APs connect over private links (LAN, WLAN, WAN or MPLS) and terminate directly on controllers. Campus APs are deployed as part of the indoor campus solution in enterprise office buildings, warehouses, hospitals, universities, and so on.

captive portal

A captive portal is a web page that allows the users to authenticate and sign in before connecting to a public-access network. Captive portals are typically used by business centers, airports, hotel lobbies, coffee shops, and other venues that offer free Wi-Fi hotspots for the guest users.

CCA

Clear Channel Assessment. In wireless networks, the CCA method detects if a channel is occupied or clear, and determines if the channel is available for data transmission.

CDP

Cisco Discovery Protocol. CDP is a proprietary Data Link Layer protocol developed by Cisco Systems. CDP runs on Cisco devices and enables networking applications to learn about the neighboring devices directly connected to the network.

CDR

Call Detail Record. A CDR contains the details of a telephone or VoIP call, such as the origin and destination addresses of the call, the start time and end time of the call, any toll charges that were added through the network or charges for operator services, and so on.

CEF

Common Event Format. The CEF is a standard for the interoperability of event or log-generating devices and applications. The standard syntax for CEF includes a prefix and a variable extension formatted as key-value pairs.

CGI

Common Gateway Interface. CGI is a standard protocol for exchanging data between the web servers and executable programs running on a server to dynamically process web pages.

CHAP

Challenge Handshake Authentication Protocol. CHAP is an authentication scheme used by PPP servers to validate the identity of remote clients.

CIDR

Classless Inter-Domain Routing. CIDR is an IP standard for creating and allocating unique identifiers for networks and devices. The CIDR IP addressing scheme is used as a replacement for the older IP addressing scheme based on classes A, B, and C. With CIDR, a single IP address can be used to designate many unique IP addresses. A CIDR IP address ends with a slash followed by the IP network prefix, for example, 192.0.2.0/24.

ClearPass

ClearPass is an access management system for creating and enforcing policies across a network to all devices and applications. The ClearPass integrated platform includes applications such as Policy Manager, Guest, Onboard, OnGuard, Insight, Profile, QuickConnect, and so on.

ClearPass Guest

ClearPass Guest is a configurable ClearPass application for secure visitor network access management.

ClearPass Policy Manager

ClearPass Policy Manager is a baseline platform for policy management, AAA, profiling, network access control, and reporting. With ClearPass Policy Manager, the network administrators can configure and manage secure network access that accommodates requirements across multiple locations and multivendor networks, regardless of device ownership and connection method.

CLI

Command-Line Interface. A console interface with a command line shell that allows users to execute text input as commands and convert these commands to appropriate functions.

CN

Common Name. CN is the primary name used to identify a certificate.

CNA

Captive Network Assistant. CNA is a popup page shown when joining a network that has a captive portal.

CoA

Change of Authorization. The RADIUS CoA is used in the AAA service framework to allow dynamic modification of the authenticated, authorized, and active subscriber sessions.

CoS

Class of Service. CoS is used in data and voice protocols for classifying packets into different types of traffic (voice, video, or data) and setting a service priority. For example, voice traffic can be assigned a higher priority over email or HTTP traffic.

CPE

Customer Premises Equipment. It refers to any terminal or equipment located at the customer premises.

CPsec

Control Plane Security. CPsec is a secure form of communication between a controller and APs to protect the control plane communications. This is performed by means of using public-key self-signed certificates created by each master controller.

CPU

Central Processing Unit. A CPU is an electronic circuitry in a computer for processing instructions.

CRC

Cyclic Redundancy Check. CRC is a data verification method for detecting errors in digital data during transmission, storage, or retrieval.

CRL

Certificate Revocation List. CRL is a list of revoked certificates maintained by a certification authority.

cryptobinding

Short for cryptographic binding. A procedure in a tunneled EAP method that binds together the tunnel protocol and the tunneled authentication methods, ensuring the relationship between a collection of data assets. Cryptographic binding focuses on protecting the server; mutual cryptographic binding protects both peer and server.

CSA

Channel Switch Announcement. The CSA element enables an AP to advertise that it is switching to a new channel before it begins transmitting on that channel. This allows the clients, which support CSA, to transition to the new channel with minimal downtime.

CSMA/CA

Carrier Sense Multiple Access / Collision Avoidance. CSMA/CA is a protocol for carrier transmission in networks using the 802.11 standard. CSMA/CA aims to prevent collisions by listening to the broadcasting nodes, and informing devices not to transmit any data until the broadcasting channel is free.

CSR

Certificate Signing Request. In PKI systems, a CSR is a message sent from an applicant to a CA to apply for a digital identity certificate.

CSV

Comma-Separated Values. A file format that stores tabular data in the plain text format separated by commas.

CTS

Clear to Send. The CTS refers to the data transmission and protection mechanism used by the 802.11 wireless networking protocol to prevent frame collision occurrences. See RTS.

CW

Contention Window. In QoS, CW refers to a window set for access categories based on the type of traffic. Based on the type and volume of the traffic, the minimum and maximum values can be calculated to provide a wider window when necessary.

DAI

Dynamic ARP inspection. A security feature that validates ARP packets in a network.

DAS

Distributed Antenna System. DAS is a network of antenna nodes strategically placed around a geographical area or structure for additional cellular coverage.

dB

Decibel. Unit of measure for sound or noise and is the difference or ratio between two signal levels.

dBm

Decibel-Milliwatts. dBm is a logarithmic measurement (integer) that is typically used in place of mW to represent receive-power level. AMP normalizes all signals to dBm, so that it is easy to evaluate performance between various vendors.

DCB

Data Center Bridging. DCB is a collection of standards developed by IEEE for creating a converged data center network using Ethernet.

DCE

Data Communication Equipment. DCE refers to the devices that establish, maintain, and terminate communication network sessions between a data source and its destination.

DCF

Distributed Coordination Function. DCF is a protocol that uses carrier sensing along with a four-way handshake to maximize the throughput while preventing packet collisions.

DDMO

Distributed Dynamic Multicast Optimization. DDMO is similar to Dynamic Multicast Optimization (DMO) where the multicast streams are converted into unicast streams on the AP instead of the controller, to enhance the quality and reliability of streaming videos, while preserving the bandwidth available to non-video clients.

DES

Data Encryption Standard. DES is a common standard for data encryption and a form of secret key cryptography, which uses only one key for encryption and decryption.

designated router

Designated router refers to a router interface that is elected to originate network link advertisements for networks using the OSPF protocol.

destination NAT

Destination Network Address Translation. Destination NAT is a process of translating the destination IP address of an end route packet in a network. Destination NAT is used for redirecting the traffic destined to a virtual host to the real host, where the virtual host is identified by the destination IP address and the real host is identified by the translated IP address.

DFS

Dynamic Frequency Selection. DFS is a mandate for radio systems operating in the 5 GHz band to be equipped with means to identify and avoid interference with Radar systems.

DFT

Discrete Fourier Transform. DFT converts discrete-time data sets into a discrete-frequency representation. See FFT.

DHCP

Dynamic Host Configuration Protocol. A network protocol that enables a server to automatically assign an IP address to an IP-enabled device from a defined range of numbers configured for a given network.

DHCP snooping

DHCP snooping enables the switch to monitor and control DHCP messages received from untrusted devices that are connected to the switch.

digital certificate

A digital certificate is an electronic document that uses a digital signature to bind a public key with an identity—information such as the name of a person or an organization, address, and so forth.

Digital wireless pulse

A wireless technology for transmitting large amounts of digital data over a wide spectrum of frequency bands with very low power for a short distance. Ultra Wideband radio can carry a huge amount of data over a distance up to 230 ft at very low power (less than 0.5 mW), and has the ability to carry signals through doors and other obstacles that tend to reflect signals at more limited bandwidths and a higher power.

Disconnect-Ack

Disconnect-Ack is a NAS response packet to a Disconnect-Request, which indicates that the session was disconnected.

Disconnect-Nak

Disconnect-Nak is NAS response packet to a Disconnect-Request, which indicates that the session was not disconnected.

Disconnect-Request

Disconnect-Request is a RADIUS packet type sent to a NAS requesting that a user or session be disconnected.

distribution certificate

Distribution certificate is used for digitally signing iOS mobile apps to enable enterprise app distribution. It verifies the identity of the app publisher.

DLNA

Digital Living Network Alliance. DLNA is a set of interoperability guidelines for sharing digital media among multimedia devices.

DMO

Dynamic Multicast Optimization. DMO is a process of converting multicast streams into unicast streams over a wireless link to enhance the quality and reliability of streaming videos, while preserving the bandwidth available to non-video clients.

DN

Distinguished Name. A series of fields in a digital certificate that, taken together, constitute the unique identity of the person or device that owns the digital certificate. Common fields in a DN include country, state, locality, organization, organizational unit, and the “common name”, which is the primary name used to identify the certificate.

DNS

Domain Name System. A DNS server functions as a phone book for the intranet and Internet users. It converts human-readable computer host names into IP addresses and IP addresses into host names. It stores several records for a domain name such as an address 'A' record, name server (NS), and mail exchanger (MX) records. The Address 'A' record is the most important record that is stored in a DNS server, because it provides the required IP address for a network peripheral or element.

DOCSIS

Data over Cable Service Interface Specification. A telecommunication standard for Internet access through cable modem.

DoS

Denial of Service. DoS is any type of attack where the attackers send excessive messages to flood traffic and thereby preventing the legitimate users from accessing the service.

DPD

Dead Peer Detection. A method used by the network devices to detect the availability of the peer devices.

DPI

Deep Packet Inspection. DPI is an advanced method of network packet filtering that is used for inspecting data packets exchanged between the devices and systems over a network. DPI functions at the Application layer of the Open Systems Interconnection (OSI) reference model and enables users to identify, categorize, track, reroute, or stop packets passing through a network.

DRT

Downloadable Regulatory Table. The DRT feature allows new regulatory approvals to be distributed for APs without a software upgrade or patch.

DS

Differentiated Services. The DS specification aims to provide uninterrupted quality of service by managing and controlling the network traffic, so that certain types of traffic get precedence.

DSCP

Differentiated Services Code Point. DSCP is a 6-bit packet header value used for traffic classification and priority assignment.

DSL

Digital Subscriber Line. The DSL technology allows the transmission of digital data over telephone lines. A DSL modem is a device used for connecting a computer or router to a telephone line that offers connectivity to the Internet.

DSSS

Direct-Sequence Spread Spectrum. DSSS is a modulation technique used for reducing overall signal interference. This technique multiplies the original data signal with a pseudo random noise spreading code. Spreading of this signal makes the resulting wideband channel more noisy, thereby increasing the resistance to interference. See FHSS.

DST

Daylight Saving Time. DST is also known as summer time that refers to the practice of advancing clocks, so that evenings have more daylight and mornings have less. Typically clocks are adjusted forward one hour near the start of spring and are adjusted backward in autumn.

DTE

Data Terminal Equipment. DTE refers to a device that converts user information into signals or re-converts the received signals.

DTIM

Delivery Traffic Indication Message. DTIM is a kind of traffic indication map. A DTIM interval determines when the APs must deliver broadcast and multicast frames to their associated clients in power save mode.

DTLS

Datagram Transport Layer Security. DTLS communications protocol provides communications security for datagram protocols.

dynamic authorization

Dynamic authorization refers to the ability to make changes to a visitor account's session while it is in progress. This might include disconnecting a session or updating some aspect of the authorization for the session.

dynamic NAT

Dynamic Network Address Translation. Dynamic NAT maps multiple public IP addresses and uses these addresses with an internal or private IP address. Dynamic NAT helps to secure a network by

masking the internal configuration of a private network.

EAP

Extensible Authentication Protocol. An authentication protocol for wireless networks that extends the methods used by the PPP, a protocol often used when connecting a computer to the Internet. EAP can support multiple authentication mechanisms, such as token cards, smart cards, certificates, one-time passwords, and public key encryption authentication.

EAP-FAST

EAP – Flexible Authentication Secure Tunnel (tunneled).

EAP-GTC

EAP – Generic Token Card. (non-tunneled).

EAP-MD5

EAP – Method Digest 5. (non-tunneled).

EAP-MSCHAP

EAP Microsoft Challenge Handshake Authentication Protocol.

EAP-MSCHAPv2

EAP Microsoft Challenge Handshake Authentication Protocol Version 2.

EAP-PEAP

EAP–Protected EAP. A widely used protocol for securely transporting authentication data across a network (tunneled).

EAP-PWD

EAP-Password. EAP-PWD is an EAP method that uses a shared password for authentication.

EAP-TLS

EAP–Transport Layer Security. EAP-TLS is a certificate-based authentication method supporting mutual authentication, integrity-protected ciphersuite negotiation and key exchange between two endpoints. See RFC 5216.

EAP-TTLS

EAP–Tunneled Transport Layer Security. EAP-TTLS is an EAP method that encapsulates a TLS session, consisting of a handshake phase and a data phase. See RFC 5281.

EAPoL

Extensible Authentication Protocol over LAN. A network port authentication protocol used in IEEE 802.1X standards to provide a generic network sign-on to access network resources.

ECC

Elliptical Curve Cryptography or Error correcting Code memory. Elliptical Curve Cryptography is a public-key encryption technique that is based on elliptic curve theory used for creating faster, smaller, and more efficient cryptographic keys. Error Correcting Code memory is a type of computer data storage that can detect and correct the most common kinds of internal data corruption. ECC

memory is used in most computers where data corruption cannot be tolerated under any circumstances, such as for scientific or financial computing.

ECDSA

Elliptic Curve Digital Signature Algorithm. ECDSA is a cryptographic algorithm that supports the use of public or private key pairs for encrypting and decrypting information.

EDCA

Enhanced Distributed Channel Access. The EDCA function in the IEEE 802.11e Quality of Service standard supports differentiated and distributed access to wireless medium based on traffic priority and Access Category types. See WMM and WME.

EIGRP

Enhanced Interior Gateway Routing Protocol. EIGRP is a routing protocol used for automating routing decisions and configuration in a network.

EIRP

Effective Isotropic Radiated Power or Equivalent Isotropic Radiated Power. EIRP refers to the output power generated when a signal is concentrated into a smaller area by the Antenna.

ESI

External Services Interface. ESI provides an open interface for integrating security solutions that solve interior network problems such as viruses, worms, spyware, and corporate compliance.

ESS

Extended Service Set. An ESS is a set of one or more interconnected BSSs that form a single sub network.

ESSID

Extended Service Set Identifier. ESSID refers to the ID used for identifying an extended service set.

Ethernet

Ethernet is a network protocol for data transmission over LAN.

EULA

End User License Agreement. EULA is a legal contract between a software application publisher or author and the users of the application.

FCC

Federal Communications Commission. FCC is a regulatory body that defines standards for the interstate and international communications by radio, television, wire, satellite, and cable.

FFT

Fast Fourier Transform. FFT is a frequency analysis mechanism that aims at faster conversion of a discrete signal in time domain into a discrete frequency domain representation. See also DFT.

FHSS

Frequency Hopping Spread Spectrum. FHSS is transmission technique that allows modulation and transmission of a data signal by rapidly switching a carrier among many frequency channels in a

random but predictable sequence. See also DSSS.

FIB

Forwarding Information Base. FIB is a forwarding table that maps MAC addresses to ports. FIB is used in network bridging, routing, and similar functions to identify the appropriate interface for forwarding packets.

FIPS

Federal Information Processing Standards. FIPS refers to a set of standards that describe document processing, encryption algorithms, and other information technology standards for use within non-military government agencies, and by government contractors and vendors who work with these agencies.

firewall

Firewall is a network security system used for preventing unauthorized access to or from a private network.

FQDN

Fully Qualified Domain Name. FQDN is a complete domain name that identifies a computer or host on the Internet.

FQLN

Fully Qualified Location Name. FQLN is a device location identifier in the format: APname.Floor.Building.Campus.

frequency allocation

Use of radio frequency spectrum as regulated by governments.

FSPL

Free Space Path Loss. FSPL refers to the loss in signal strength of an electromagnetic wave that would result from a line-of-sight path through free space (usually air), with no obstacles nearby to cause reflection or diffraction.

FTP

File Transfer Protocol. A standard network protocol used for transferring files between a client and server on a computer network.

GARP

Generic Attribute Registration Protocol. GVRP is a LAN protocol that allows the network nodes to register and de-register attributes, such as network addresses, with each other.

GAS

Generic Advertisement Service. GAS is a request-response protocol, which provides Layer 2 transport mechanism between a wireless client and a server in the network prior to authentication. It helps in determining a wireless network infrastructure before associating clients, and allows clients to send queries to multiple 802.11 networks in parallel.

gateway

Gateway is a network node that allows traffic to flow in and out of the network.

Gbps

Gigabits per second.

GBps

Gigabytes per second.

GET

GET refers HTTP request method or an SNMP operation method. The GET HTTP request method submits data to be processed to a specified resource. The GET SNMP operation method obtains information from the Management Information Base (MIB).

GHz

Gigahertz.

GMT

Greenwich Mean Time. GMT refers to the mean solar time at the Royal Observatory in Greenwich, London. GMT is the same as Coordinated Universal Time (UTC) standard, written as an offset of UTC +/- 00:00.

goodput

Goodput is the application level throughput that refers to the ratio of the total bytes transmitted or received in the network to the total air time required for transmitting or receiving the bytes.

GPS

Global Positioning System. A satellite-based global navigation system.

GRE

Generic Routing Encapsulation. GRE is an IP encapsulation protocol that is used to transport packets over a network.

GTC

Generic Token Card. GTC is a protocol that can be used as an alternative to MSCHAPv2 protocol. GTC allows authentication to various authentication databases even in cases where MSCHAPv2 is not supported by the database.

GVRP

GARP VLAN Registration Protocol or Generic VLAN Registration Protocol. GARP is an IEEE 802.1Q-compliant protocol that facilitates VLAN registration and controls VLANs within a larger network.

H2QP

Hotspot 2.0 Query Protocol.

hot zone

Wireless access area created by multiple hotspots that are located in close proximity to one another. Hot zones usually combine public safety APs with public hotspots.

hotspot

Hotspot refers to a WLAN node that provides Internet connection and virtual private network (VPN) access from a given location. A business traveler, for example, with a laptop equipped for Wi-Fi can

look up a local hotspot, contact it, and get connected through its network to reach the Internet.

HSPA

High-Speed Packet Access.

HT

High Throughput. IEEE 802.11n is an HT WLAN standard that aims to achieve physical data rates of close to 600 Mbps on the 2.4 GHz and 5 GHz bands.

HTTP

Hypertext Transfer Protocol. The HTTP is an application protocol to transfer data over the web. The HTTP protocol defines how messages are formatted and transmitted, and the actions that the w servers and browsers should take in response to various commands.

HTTPS

Hypertext Transfer Protocol Secure. HTTPS is a variant of the HTTP that adds a layer of security on the data in transit through a secure socket layer or transport layer security protocol connection.

IAS

Internet Authentication Service. IAS is a component of Windows Server operating systems that provides centralized user authentication, authorization, and accounting.

ICMP

Internet Control Message Protocol. ICMP is an error reporting protocol. It is used by network devices such as routers, to send error messages and operational information to the source IP address when network problems prevent delivery of IP packets.

IDS

Intrusion Detection System. IDS monitors a network or systems for malicious activity or policy violations and reports its findings to the management system deployed in the network.

IEEE

Institute of Electrical and Electronics Engineers.

IGMP

Internet Group Management Protocol. Communications protocol used by hosts and adjacent routers on IP networks to establish multicast group memberships.

IGMP snooping

IGMP snooping prevents multicast flooding on Layer 2 network by treating multicast traffic as broadcast traffic. Without IGMP snooping, all streams could be flooded to all ports on that VLAN. When multicast flooding occurs, end-hosts that happen to be in the same VLAN would receive all the streams only to be discarded without snooping.

IGP

Interior Gateway Protocol. IGP is used for exchanging routing information between gateways within an autonomous system (for example, a system of corporate local area networks).

IGRP

Interior Gateway Routing Protocol. IGRP is a distance vector interior routing protocol used by routers to exchange routing data within an autonomous system.

IKE

Internet Key Exchange. IKE is a key management protocol used with IPsec protocol to establish a secure communication channel. IKE provides additional feature, flexibility, and ease of configuration for IPsec standard.

IKEv1

Internet Key Exchange version 1. IKEv1 establishes a secure authenticated communication channel by using either the pre-shared key (shared secret), digital signatures, or public key encryption. IKEv1 operates in Main and Aggressive modes. See RFC 2409.

IKEv2

Internet Key Exchange version 2. IKEv2 uses the secure channel established in Phase 1 to negotiate Security Associations on behalf of services such as IPsec. IKEv2 uses pre-shared key and Digital Signature for authentication. See RFC 4306.

IoT

Internet of Things. IoT refers to the internetworking of devices that are embedded with electronics, software, sensors, and network connectivity features allowing data exchange over the Internet.

IPM

Intelligent Power Monitoring. IPM is a feature supported on certain APs that actively measures the power utilization of an AP and dynamically adapts to the power resources.

IPS

Intrusion Prevention System. The IPS monitors a network for malicious activities such as security threats or policy violations. The main function of an IPS is to identify suspicious activity, log the information, attempt to block the activity, and report it.

IPsec

Internet Protocol security. IPsec is a protocol suite for secure IP communications that authenticates and encrypts each IP packet in a communication session.

IPSG

Internet Protocol Source Guard. IPSG restricts IP address from untrusted interface by filtering traffic based on list of addresses in the DHCP binding database or manually configured IP source bindings. It prevents IP spoofing attacks.

IrDA

An industry-sponsored organization set up in 1993 to create international standards for the hardware and software used in infrared communication links. In this special form of radio transmission, a focused ray of light in the infrared frequency spectrum, measured in terahertz (THz), or trillions of hertz (cycles per second), is modulated with information and sent from a transmitter to a receiver over a relatively short distance.

ISAKMP

Internet Security Association and Key Management Protocol. ISAKMP is used for establishing Security Associations and cryptographic keys in an Internet environment.

ISP

Internet Service Provider. An ISP is an organization that provides services for accessing and using the Internet.

JSON

JavaScript Object Notation. JSON is an open-standard, language-independent, lightweight data-interchange format used to transmit data objects consisting of attribute–value pairs. JSON uses a "self-describing" text format that is easy for humans to read and write, and that can be used as a data format by any programming language.

Kbps

Kilobits per second.

KBps

Kilobytes per second.

keepalive

Signal sent at periodic intervals from one device to another to verify that the link between the two devices is working. If no reply is received, data will be sent by a different path until the link is restored. A keepalive can also be used to indicate that the connection should be preserved so that the receiving device does not consider it timed out and drop it.

L2TP

Layer-2 Tunneling Protocol. L2TP is a networking protocol used by the ISPs to enable VPN operations.

LACP

Link Aggregation Control Protocol. LACP is used for the collective handling of multiple physical ports that can be seen as a single channel for network traffic purposes.

LAG

Link Aggregation Group . A LAG combines a number of physical ports together to make a single high-bandwidth data path. LAGs can connect two switches to provide a higher-bandwidth connection to a public network.

LAN

Local Area Network. A LAN is a network of connected devices within a distinct geographic area such as an office or a commercial establishment and share a common communications line or wireless link to a server.

LCD

Liquid Crystal Display. LCD is the technology used for displays in notebook and other smaller computers. Like LED and gas-plasma technologies, LCDs allow displays to be much thinner than the cathode ray tube technology.

LDAP

Lightweight Directory Access Protocol. LDAP is a communication protocol that provides the ability to access and maintain distributed directory information services over a network.

LDPC

Low-Density Parity-Check. LDPC is a method of transmitting a message over a noisy transmission channel using a linear error correcting code. An LDPC is constructed using a sparse bipartite graph.

LEAP

Lightweight Extensible Authentication Protocol. LEAP is a Cisco proprietary version of EAP used in wireless networks and Point-to-Point connections.

LED

Light Emitting Diode. LED is a semiconductor light source that emits light when an electric current passes through it.

LEEF

Log Event Extended Format. LEEF is a type of customizable syslog event format. An extended log file contains a sequence of lines containing ASCII characters terminated by either the sequence LF or CRLF.

LI

Lawful Interception. LI refers to the procedure of obtaining communications network data by the Law Enforcement Agencies for the purpose of analysis or evidence.

LLDP

Link Layer Discovery Protocol. LLDP is a vendor-neutral link layer protocol in the Internet Protocol suite used by network devices for advertising their identity, capabilities, and neighbors on an IEEE 802 local area network, which is principally a wired Ethernet.

LLDP-MED

LLDP–Media Endpoint Discovery. LLDP-MED facilitates information sharing between endpoints and network infrastructure devices.

LMS

Local Management Switch. In multi-controller networks, each controller acts as an LMS and terminates user traffic from the APs, processes, and forwards the traffic to the wired network.

LNS

L2TP Network Server. LNS is an equipment that connects to a carrier and handles the sessions from broadband lines. It is also used for dial-up and mobile links. LNS handles authentication and routing of the IP addresses. It also handles the negotiation of the link with the equipment and establishes a session.

LTE

Long Term Evolution. LTE is a 4G wireless communication standard that provides high-speed wireless communication for mobile phones and data terminals. See 4G.

MAB

MAC Authentication Bypass. Endpoints such as network printers, Ethernet-based sensors, cameras, and wireless phones do not support 802.1X authentication. For such endpoints, MAC Authentication Bypass mechanism is used. In this method, the MAC address of the endpoint is used to authenticate the endpoint.

MAC

Media Access Control. A MAC address is a unique identifier assigned to network interfaces for communications on a network.

MAM

Mobile Application Management. MAM refers to software and services used to secure, manage, and distribute mobile applications used in enterprise settings on mobile devices like smartphones and tablet computers. Mobile Application Management can apply to company-owned mobile devices as well as BYOD.

Mbps

Megabits per second

MBps

Megabytes per second

MCS

Modulation and Coding Scheme. MCS is used as a parameter to determine the data rate of a wireless connection for high throughput.

MD4

Message Digest 4. MD4 is an earlier version of MD5 and is an algorithm used to verify data integrity through the creation of a 128-bit message digest from data input.

MD5

Message Digest 5. The MD5 algorithm is a widely used hash function producing a 128-bit hash value from the data input.

MDAC

Microsoft Data Access Components. MDAC is a framework of interrelated Microsoft technologies that provides a standard database for Windows OS.

MDM

Mobile Device Management. MDM is an administrative software to manage, monitor, and secure mobile devices of the employees in a network.

mDNS

Multicast Domain Name System. mDNS provides the ability to perform DNS-like operations on the local link in the absence of any conventional unicast DNS server. The mDNS protocol uses IP multicast User Datagram Protocol (UDP) packets, and is implemented by the Apple Bonjour and Linux NSS-mDNS services. mDNS works in conjunction with DNS Service Discovery (DNS-SD), a companion zero-configuration technique specified. See RFC 6763.

MFA

Multi-factor Authentication. MFA lets you require multiple factors, or proofs of identity, when authenticating a user. Policy configurations define how often multi-factor authentication will be required, or conditions that will trigger it.

MHz

Megahertz

MIB

Management Information Base. A hierarchical database used by SNMP to manage the devices being monitored.

microwave

Electromagnetic energy with a frequency higher than 1 GHz, corresponding to wavelength shorter than 30 centimeters.

MIMO

Multiple Input Multiple Output. An antenna technology for wireless communications in which multiple antennas are used at both source (transmitter) and destination (receiver). The antennas at each end of the communications circuit are combined to minimize errors and optimize data speed.

MISO

Multiple Input Single Output. An antenna technology for wireless communications in which multiple antennas are used at the source (transmitter). The antennas are combined to minimize errors and optimize data speed. The destination (receiver) has only one antenna.

MLD

Multicast Listener Discovery. A component of the IPv6 suite. It is used by IPv6 routers for discovering multicast listeners on a directly attached link.

MPDU

MAC Protocol Data Unit. MPDU is a message exchanged between MAC entities in a communication system based on the layered OSI model.

MPLS

Multiprotocol Label Switching. The MPLS protocol speeds up and shapes network traffic flows.

MPPE

Microsoft Point-to-Point Encryption. A method of encrypting data transferred across PPP-based dial-up connections or PPTP-based VPN connections.

MS-CHAP

Microsoft Challenge Handshake Authentication Protocol. MS-CHAP is Password-based, challenge-response, mutual authentication protocol that uses MD4 and DES encryption.

MS-CHAPv1

Microsoft Challenge Handshake Authentication Protocol version 1. MS-CHAPv1 extends the user authentication functionality provided on Windows networks to remote workstations. MS-CHAPv1 supports only one-way authentication.

MS-CHAPv2

Microsoft Challenge Handshake Authentication Protocol version 2. MS-CHAPv2 is an enhanced version of the MS-CHAP protocol that supports mutual authentication.

MSS

Maximum Segment Size. MSS is a parameter of the options field in the TCP header that specifies the largest amount of data, specified in bytes, that a computer or communications device can receive in a single TCP segment.

MSSID

Mesh Service Set Identifier. MSSID is the SSID used by the client to access a wireless mesh network.

MSTP

Multiple Spanning Tree Protocol. MSTP configures a separate Spanning Tree for each VLAN group and blocks all but one of the possible alternate paths within each spanning tree.

MTU

Maximum Transmission Unit. MTU is the largest size packet or frame specified in octets (eight-bit bytes) that can be sent in networks such as the Internet.

MU-MIMO

Multi-User Multiple-Input Multiple-Output. MU-MIMO is a set of multiple-input and multiple-output technologies for wireless communication, in which users or wireless terminals with one or more antennas communicate with each other.

MVRP

Multiple VLAN Registration Protocol. MVRP is a Layer 2 network protocol used for automatic configuration of VLAN information on switches.

mW

milliWatts. mW is 1/1000 of a Watt. It is a linear measurement (always positive) that is generally used to represent transmission.

NAC

Network Access Control. NAC is a computer networking solution that uses a set of protocols to define and implement a policy that describes how devices can secure access to network nodes when they initially attempt to connect to a network.

NAD

Network Access Device. NAD is a device that automatically connects the user to the preferred network, for example, an AP or an Ethernet switch.

NAK

Negative Acknowledgement. NAK is a response indicating that a transmitted message was received with errors or it was corrupted, or that the receiving end is not ready to accept transmissions.

NAP

Network Access Protection. The NAP feature in the Windows Server allows network administrators to define specific levels of network access based on identity, groups, and policy compliance. The NAP

Agent is a service that collects and manages health information for NAP client computers. If a client is not compliant, NAP provides a mechanism to automatically bring the client back into compliance and then dynamically increase its level of network access.

NAS

Network Access Server. NAS provides network access to users, such as a wireless AP, network switch, or dial-in terminal server.

NAT

Network Address Translation. NAT is a method of remapping one IP address space into another by modifying network address information in Internet Protocol (IP) datagram packet headers while they are in transit across a traffic routing device.

NetBIOS

Network Basic Input/Output System. A program that lets applications on different computers communicate within a LAN.

netmask

Netmask is a 32-bit mask used for segregating IP address into subnets. Netmask defines the class and range of IP addresses.

NFC

Near-Field Communication. NFC is a short-range wireless connectivity standard (ECMA-340, ISO/IEC 18092) that uses magnetic field induction to enable communication between devices when they touch or are brought closer (within a few centimeters of distance). The standard specifies a way for the devices to establish a peer-to-peer (P2P) network to exchange data.

NIC

Network Interface Card. NIC is a hardware component that allows a device to connect to the network.

Nmap

Network Mapper. Nmap is an open-source utility for network discovery and security auditing. Nmap uses IP packets to determine such things as the hosts available on a network and their services, operating systems and versions, types of packet filters/firewalls, and so on.

NMI

Non-Maskable Interrupt. NMI is a hardware interrupt that standard interrupt-masking techniques in the system cannot ignore. It typically occurs to signal attention for non-recoverable hardware errors.

NMS

Network Management System. NMS is a set of hardware and/or software tools that allow an IT professional to supervise the individual components of a network within a larger network management framework.

NOE

New Office Environment. NOE is a proprietary VoIP protocol designed by Alcatel-Lucent Enterprise.

NTP

Network Time Protocol. NTP is a protocol for synchronizing the clocks of computers over a network.

OAuth

Open Standard for Authorization. OAuth is a token-based authorization standard that allows websites or third-party applications to access user information, without exposing the user credentials.

OCSP

Online Certificate Status Protocol. OCSP is used for determining the current status of a digital certificate without requiring a CRL.

OFDM

Orthogonal Frequency Division Multiplexing. OFDM is a scheme for encoding digital data on multiple carrier frequencies.

OID

Object Identifier. An OID is an identifier used to name an object. The OIDs represent nodes or managed objects in a MIB hierarchy. The OIDs are designated by text strings and integer sequences and are formally defined as per the ASN.1 standard.

OKC

Opportunistic Key Caching. OKC is a technique available for authentication between multiple APs in a network where those APs are under common administrative control. Using OKC, a station roaming to any AP in the network will not have to complete a full authentication exchange, but will instead just perform the 4-way handshake to establish transient encryption keys.

onboarding

The process of preparing a device for use on an enterprise network, by creating the appropriate access credentials and setting up the network connection parameters.

OpenFlow

OpenFlow is an open communications interface between control plane and the forwarding layers of a network.

OpenFlow agent

OpenFlow agent. OpenFlow is a software module in Software-Defined Networking (SDN) that allows the abstraction of any legacy network element, so that it can be integrated and managed by the SDN controller. OpenFlow runs on network devices such as switches, routers, wireless controllers, and APs.

Optical wireless

Optical wireless is combined use of conventional radio frequency wireless and optical fiber for telecommunication. Long-range links are provided by using optical fibers; the links from the long-range endpoints to end users are accomplished by RF wireless or laser systems. RF wireless at Ultra High Frequencies and microwave frequencies can carry broadband signals to individual computers at substantial data speeds.

OSI

Open Systems Interconnection. OSI is a reference model that defines a framework for communication between the applications in a network.

OSPF

Open Shortest Path First. OSPF is a link-state routing protocol for IP networks. It uses a link-state routing algorithm and falls into the group of interior routing protocols that operates within a single Autonomous System (AS).

OSPFv2

Open Shortest Path First version 2. OSPFv2 is the version 2 of the link-state routing protocol, OSPF. See RFC 2328.

OUI

Organizationally Unique Identifier. Synonymous with company ID or vendor ID, an OUI is a 24-bit, globally unique assigned number, referenced by various standards. The first half of a MAC address is OUI.

OVA

Open Virtualization Archive. OVA contains a compressed installable version of a virtual machine.

OVF

Open Virtualization Format. OVF is a specification that describes an open-standard, secure, efficient, portable and extensible format for packaging and distributing software for virtual machines.

PAC

Protected Access Credential. PAC is distributed to clients for optimized network authentication. These credentials are used for establishing an authentication tunnel between the client and the authentication server.

PAP

Password Authentication Protocol. PAP validates users by password. PAP does not encrypt passwords for transmission and is thus considered insecure.

PAPI

Process Application Programming Interface. PAPI controls channels for ARM and Wireless Intrusion Detection System (WIDS) communication to the master controller. A separate PAPI control channel connects to the local controller where the SSID tunnels terminate.

PBR

Policy-based Routing. PBR provides a flexible mechanism for forwarding data packets based on polices configured by a network administrator.

PDU

Power Distribution Unit or Protocol Data Unit. Power Distribution Unit is a device that distributes electric power to the networking equipment located within a data center. Protocol Data Unit contains protocol control Information that is delivered as a unit among peer entities of a network.

PEAP

Protected Extensible Authentication Protocol. PEAP is a type of EAP communication that addresses security issues associated with clear text EAP transmissions by creating a secure channel encrypted and protected by TLS.

PEF

Policy Enforcement Firewall. PEF also known as PEFNG provides context-based controls to enforce application-layer security and prioritization. The customers using Aruba mobility controllers can avail PEF features and services by obtaining a PEF license. PEF for VPN users—Customers with PEF for VPN license can apply firewall policies to the user traffic routed to a controller through a VPN tunnel.

PEFNG

Policy Enforcement Firewall. PEF also known as PEFNG provides context-based controls to enforce application-layer security and prioritization. The customers using Aruba mobility controllers can avail PEF features and services by obtaining a PEF license. PEF for VPN users—Customers with PEF for VPN license can apply firewall policies to the user traffic routed to a controller through a VPN tunnel.

PEFV

Policy Enforcement Firewall. PEF also known as PEFNG provides context-based controls to enforce application-layer security and prioritization. The customers using Aruba mobility controllers can avail PEF features and services by obtaining a PEF license. PEF for VPN users—Customers with PEF for VPN license can apply firewall policies to the user traffic routed to a controller through a VPN tunnel.

PFS

Perfect Forward Secrecy. PFS refers to the condition in which a current session key or long-term private key does not compromise the past or subsequent keys.

PHB

Per-hop behavior. PHB is a term used in DS or MPLS. It defines the policy and priority applied to a packet when traversing a hop (such as a router) in a DiffServ network.

PIM

Protocol-Independent Multicast. PIM refers to a family of multicast routing protocols for IP networks that provide one-to-many and many-to-many distribution of data over a LAN, WAN, or the Internet.

PIN

Personal Identification Number. PIN is a numeric password used to authenticate a user to a system.

PKCS#n

Public-key cryptography standard n. PKCS#n refers to a numbered standard related to topics in cryptography, including private keys (PKCS#1), digital certificates (PKCS#7), certificate signing requests (PKCS#10), and secure storage of keys and certificates (PKCS#12).

PKI

Public Key Infrastructure. PKI is a security technology based on digital certificates and the assurances provided by strong cryptography. See also certificate authority, digital certificate, public key, private key.

PLMN

Public Land Mobile Network. PLMS is a network established and operated by an administration or by a Recognized Operating Agency for the specific purpose of providing land mobile telecommunications services to the public.

PMK

Pairwise Master Key. PMK is a shared secret key that is generated after PSK or 802.1X authentication.

PoE

Power over Ethernet. PoE is a technology for wired Ethernet LANs to carry electric power required for the device in the data cables. The IEEE 802.3af PoE standard provides up to 15.4 W of power on each port.

PoE+

Power over Ethernet+. PoE+ is an IEEE 802.3at standard that provides 25.5W power on each port.

POST

Power On Self Test. An HTTP request method that requests data from a specified resource.

PPP

Point-to-Point Protocol. PPP is a data link (layer 2) protocol used to establish a direct connection between two nodes. It can provide connection authentication, transmission encryption, and compression.

PPPoE

Point-to-Point Protocol over Ethernet. PPPoE is a method of connecting to the Internet, typically used with DSL services, where the client connects to the DSL modem.

PPTP

Point-to-Point Tunneling Protocol. PPTP is a method for implementing virtual private networks. It uses a control channel over TCP and a GRE tunnel operating to encapsulate PPP packets.

private key

The part of a public-private key pair that is always kept private. The private key encrypts the signature of a message to authenticate the sender. The private key also decrypts a message that was encrypted with the public key of the sender.

PRNG

Pseudo-Random Number Generator. PRNG is an algorithm for generating a sequence of numbers whose properties approximate the properties of sequences of random numbers.

PSK

Pre-shared key. A unique shared secret that was previously shared between two parties by using a secure channel. This is used with WPA security, which requires the owner of a network to provide a passphrase to users for network access.

PSU

Power Supply Unit. PSU is a unit that supplies power to an equipment by converting mains AC to low-voltage regulated DC power.

public key

The part of a public-private key pair that is made public. The public key encrypts a message and the message is decrypted with the private key of the recipient.

PVST

Per-VLAN Spanning Tree. PVST provides load balancing of VLANs across multiple ports resulting in optimal usage of network resources.

PVST+

Per-VLAN Spanning Tree+. PVST+ is an extension of the PVST standard that uses the 802.1Q trunking technology.

QoS

Quality of Service. It refers to the capability of a network to provide better service and performance to a specific network traffic over various technologies.

RA

Router Advertisement. The RA messages are sent by the routers in the network when the hosts send multicast router solicitation to the multicast address of all routers.

Radar

Radio Detection and Ranging. Radar is an object-detection system that uses radio waves to determine the range, angle, or velocity of objects.

RADIUS

Remote Authentication Dial-In User Service. An Industry-standard network access protocol for remote authentication. It allows authentication, authorization, and accounting of remote users who want to access network resources.

RAM

Random Access Memory.

RAPIDS

Rogue Access Point identification and Detection System. An AMP module that is designed to identify and locate wireless threats by making use of all of the information available from your existing infrastructure.

RARP

Reverse Address Resolution Protocol. RARP is a protocol used by a physical machine in a local area network for determining the IP address from the ARP table or cache of the gateway server.

Regex

Regular Expression. Regex refers to a sequence of symbols and characters defining a search pattern.

Registration Authority

Type of Certificate Authority that processes certificate requests. The Registration Authority verifies that requests are valid and comply with certificate policy, and authenticates the user's identity. The Registration Authority then forwards the request to the Certificate Authority to sign and issue the certificate.

Remote AP

Remote APs extend corporate network to the users working from home or at temporary work sites. Remote APs are deplyed at branch office sites and are connected to the central network on a WAN link.

REST

Representational State Transfer. REST is a simple and stateless architecture that the web services use for providing interoperability between computer systems on the Internet. In a RESTful web service, requests made to the URI of a resource will elicit a response that may be in XML, HTML, JSON or some other defined format.

RF

Radio Frequency. RF refers to the electromagnetic wave frequencies within a range of 3 kHz to 300 GHz, including the frequencies used for communications or Radar signals.

RFC

Request For Comments. RFC is a commonly used format for the Internet standards documentss.

RFID

Radio Frequency Identification. RFID uses radio waves to automatically identify and track the information stored on a tag attached to an object.

RIP

Routing Information Protocol. RIP prevents the routing loops by limiting the number of hops allowed in a path from source to destination.

RJ45

Registered Jack 45. RJ45 is a physical connector for network cables.

RMA

Return Merchandise Authorization. RMA is a part of the product returning process that authorizes users to return a product to the manufacturer or distributor for a refund, replacement, or repair. The customers who want to return a product within its Warranty period contact the manufacturer to initiate the product returning process. The manufacturer or the seller generates an authorization number for the RMA, which is used by the customers, when returning a product to the warehouse.

RMON

Remote Monitoring. RMON provides standard information that a network administrator can use to monitor, analyze, and troubleshoot a group of distributed LANs.

RoW

Rest of World. RoW or RW is an operating country code of a device.

RSA

Rivest, Shamir, Adleman. RSA is a cryptosystem for public-key encryption, and is widely used for securing sensitive data, particularly when being sent over an insecure network such as the Internet.

RSSI

Received Signal Strength Indicator. RSSI is a mechanism by which RF energy is measured by the circuitry on a wireless NIC (0-255). The RSSI is not standard across vendors. Each vendor determines its own RSSI scale/values.

RSTP

Rapid Spanning Tree Protocol. RSTP provides significantly faster spanning tree convergence after a topology change, introducing new convergence behaviors and bridge port roles to do this.

RTCP

RTP Control Protocol. RTCP provides out-of-band statistics and control information for an Real-Time Transport Protocol session.

RTLS

Real-Time Location Systems. RTLS automatically identifies and tracks the location of objects or people in real time, usually within a building or other contained area.

RTP

Real-Time Transport Protocol. RTP is a network protocol used for delivering audio and video over IP networks.

RTS

Request to Send. RTS refers to the data transmission and protection mechanism used by the 802.11 wireless networking protocol to prevent frame collision occurrences. See CTS.

RTSP

Real Time Streaming Protocol. RTSP is a network control protocol designed for use in entertainment and communications systems to control streaming media servers.

RVI

Routed VLAN Interface. RVI is a switch interface that forwards packets between VLANs.

RW

Rest of World. RoW or RW is an operating country code of a device.

SA

Security Association. SA is the establishment of shared security attributes between two network entities to support secure communication.

SAML

Security Assertion Markup Language. SAML is an XML-based framework for communicating user authentication, entitlement, and attribute information. SAML enables single sign-on by allowing users to authenticate at an identity provider and then access service providers without additional authentication.

SCEP

Simple Certificate Enrollment Protocol. SCEP is a protocol for requesting and managing digital certificates.

SCP

Secure Copy Protocol. SCP is a network protocol that supports file transfers between hosts on a network.

SCSI

Small Computer System Interface. SCSI refers to a set of interface standards for physical connection and data transfer between a computer and the peripheral devices such as printers, disk drives, CD-ROM, and so on.

SD-WAN

Software-Defined Wide Area Network. SD-WAN is an application for applying SDN technology to WAN connections that connect enterprise networks across disparate geographical locations.

SDN

Software-Defined Networking. SDN is an umbrella term encompassing several kinds of network technology aimed at making the network as agile and flexible as the virtualized server and storage infrastructure of the modern data center.

SDR

Server Derivation Rule. An SDR refers to a role assignment model used by the controllers running ArubaOS to assign roles and VLANs to the WLAN users based on the rules defined under a server group. The SDRs override the default authentication roles and VLANs defined in the AAA and Virtual AP profiles.

SDU

Service Data Unit. SDU is a unit of data that has been passed down from an OSI layer to a lower layer and that has not yet been encapsulated into a PDU by the lower layer.

SFP

The Small Form-factor Pluggable. SFP is a compact, hot-pluggable transceiver that is used for both telecommunication and data communications applications.

SFP+

Small Form-factor Pluggable+. SFP+ supports up to data rates up to 16 Gbps.

SFTP

Secure File Transfer Protocol. SFTP is a network protocol that allows file access, file transfer, and file management functions over a secure connection.

SHA

Secure Hash Algorithm. SHA is a family of cryptographic hash functions. The SHA algorithm includes the SHA, SHA-1, SHA-2 and SHA-3 variants.

SIM

Subscriber Identity Module. SIM is an integrated circuit that is intended to securely store the International Mobile Subscriber Identity (IMSI) number and its related key, which are used for identifying and authenticating subscribers on mobile telephony devices.

SIP

Session Initiation Protocol. SIP is used for signaling and controlling multimedia communication session such as voice and video calls.

SIRT

Security Incident Response Team. SIRT is responsible for reviewing as well as responding to computer security incident reports and activity.

SKU

Stock Keeping Unit. SKU refers to the product and service identification code for the products in the inventory.

SLAAC

Stateless Address Autoconfiguration. SLAAC provides the ability to address a host based on a network prefix that is advertised from a local network router through router advertisements.

SMB

Server Message Block or Small and Medium Business. Server Message Block operates as an application-layer network protocol mainly used for providing shared access to files, printers, serial ports, and for miscellaneous communications between the nodes on a network.

SMS

Short Message Service. SMS refers to short text messages (up to 140 characters) sent and received through mobile phones.

SMTP

Simple Mail Transfer Protocol. SMTP is an Internet standard protocol for electronic mail transmission.

SNIR

Signal-to-Noise-Plus-Interference Ratio. SNIR refers to the power of a central signal of interest divided by the sum of the interference power and the power of the background noise. SINR is defined as the power of a certain signal of interest divided by the sum of the interference power (from all the other interfering signals) and the power of some background noise.

SNMP

Simple Network Management Protocol. SNMP is a TCP/IP standard protocol for managing devices on IP networks. Devices that typically support SNMP include routers, switches, servers, workstations, printers, modem racks, and more. It is used mostly in network management systems to monitor network-attached devices for conditions that warrant administrative attention.

SNMPv1

Simple Network Management Protocol version 1. SNMPv1 is a widely used network management protocol.

SNMPv2

Simple Network Management Protocol version 2. SNMPv2 is an enhanced version of SNMPv1, which includes improvements in the areas of performance, security, confidentiality, and manager-to-manager communications.

SNMPv2c

Community-Based Simple Network Management Protocol version 2. SNMPv2C uses the community-based security scheme of SNMPv1 and does not include the SNMPv2 security model.

SNMPv3

Simple Network Management Protocol version 3. SNMPv3 is an enhanced version of SNMP that includes security and remote configuration features.

SNR

Signal-to-Noise Ratio. SNR is used for comparing the level of a desired signal with the level of background noise.

SNTP

Simple Network Time Protocol. SNTP is a less complex implementation of NTP. It uses the same , but does not require the storage of state over extended periods of time.

SOAP

Simple Object Access Protocol. SOAP enables communication between the applications running on different operating systems, with different technologies and programming languages. SOAP is an XML-based messaging protocol for exchanging structured information between the systems that support web services.

SoC

System on a Chip. SoC is an Integrated Circuit that integrates all components of a computer or other electronic system into a single chip.

source NAT

Source NAT changes the source address of the packets passing through the router. Source NAT is typically used when an internal (private) host initiates a session to an external (public) host.

SSH

Secure Shell. SSH is a network protocol that provides secure access to a remote device.

SSID

Service Set Identifier. SSID is a name given to a WLAN and is used by the client to access a WLAN network.

SSL

Secure Sockets Layer. SSL is a computer networking protocol for securing connections between network application clients and servers over the Internet.

SSO

Single Sign-On. SSO is an access-control property that allows the users to log in once to access multiple related, but independent applications or systems to which they have privileges. The process authenticates the user across all allowed resources during their session, eliminating additional login prompts.

STBC

Space-Time Block Coding. STBC is a technique used in wireless communications to transmit multiple copies of a data stream across a number of antennas and to exploit the various received versions of the data to improve the reliability of data transfer.

STM

Station Management. STM is a process that handles AP management and user association.

STP

Spanning Tree Protocol. STP is a network protocol that builds a logical loop-free topology for Ethernet networks.

SU-MIMO

Single-User Multiple-Input Multiple-Output. SU-MIMO allocates the full bandwidth of the AP to a single high-speed device during the allotted time slice.

subnet

Subnet is the logical division of an IP network.

subscription

A business model where a customer pays a certain amount as subscription price to obtain access to a product or service.

SVP

SpectraLink Voice Priority. SVP is an open, straightforward QoS approach that has been adopted by most leading vendors of WLAN APs. SVP favors isochronous voice packets over asynchronous data packets when contending for the wireless medium and when transmitting packets onto the wired LAN.

SWAN

Structured Wireless-Aware Network. A technology that incorporates a Wireless Local Area Network (WLAN) into a wired Wide Area Network (WAN). SWAN technology can enable an existing wired network to serve hundreds of users, organizations, corporations, or agencies over a large geographic area. SWAN is said to be scalable, secure, and reliable.

TAC

Technical Assistance Center.

TACACS

Terminal Access Controller Access Control System. TACACS is a family of protocols that handles remote authentication and related services for network access control through a centralized server.

TACACS+

Terminal Access Controller Access Control System+. TACACS+ provides separate authentication, authorization, and accounting services. It is derived from, but not backward compatible with, TACACS.

TCP

Transmission Control Protocol. TCP is a communication protocol that defines the standards for establishing and maintaining network connection for applications to exchange data.

TCP/IP

Transmission Control Protocol/ Internet Protocol. TCP/IP is the basic communication language or protocol of the Internet.

TEAP

Tunnel Extensible Authentication Protocol (TEAP) is a tunnel-based EAP method that enables secure communication between a peer and a server by using the Transport Layer Security (TLS) protocol to establish a mutually authenticated tunnel. Within the tunnel, TLV objects are used to convey authentication-related data between the EAP peer and the EAP server.

TFTP

Trivial File Transfer Protocol. The TFTP is a software utility for transferring files from or to a remote host.

TIM

Traffic Indication Map. TIM is an information element that advertises if any associated stations have buffered unicast frames. APs periodically send the TIM within a beacon to identify the stations that are using power saving mode and the stations that have undelivered data buffered on the AP.

TKIP

Temporal Key Integrity Protocol. A part of the WPA encryption standard for wireless networks. TKIP is the next-generation Wired Equivalent Privacy (WEP) that provides per-packet key mixing to address the flaws encountered in the WEP standard.

TLS

Transport Layer Security. TLS is a cryptographic protocol that provides communication security over the Internet. TLS encrypts the segments of network connections above the Transport Layer by using asymmetric cryptography for key exchange, symmetric encryption for privacy, and message authentication codes for message integrity.

TLV

Type-length-value or Tag-Length-Value. TLV is an encoding format. It refers to the type of data being processed, the length of the value, and the value for the type of data being processed.

ToS

Type of Service. The ToS field is part of the IPv4 header, which specifies datagrams priority and requests a route for low-delay, high-throughput, or a highly reliable service.

TPC

Transmit Power Control. TPC is a part of the 802.11h amendment. It is used to regulate the power levels used by 802.11a radio cards.

TPM

Trusted Platform Module. TPM is an international standard for a secure cryptoprocessor, which is a dedicated microcontroller designed to secure hardware by integrating cryptographic keys into

devices.

TSF

Timing Synchronization Function. TSF is a WLAN function that is used for synchronizing the timers for all the stations in a BSS.

TSPEC

Traffic Specification. TSPEC allows an 802.11e client or a QoS-capable wireless client to signal its traffic requirements to the AP.

TSV

Tab-Separated Values. TSV is a file format that allows the exchange of tabular data between applications that use different internal data formats.

TTL

Time to Live. TTL or hop limit is a mechanism that sets limits for data expiry in a computer or network.

TTY

TeleTypeWriter. TTY-enabled devices allow telephones to transmit text communications for people who are deaf or hard of hearing as well as transmit voice communication.

TXOP

Transmission Opportunity. TXOP is used in wireless networks supporting the IEEE 802.11e Quality of Service (QoS) standard. Used in both EDCA and HCF Controlled Channel Access modes of operation, TXOP is a bounded time interval in which stations supporting QoS are permitted to transfer a series of frames. TXOP is defined by a start time and a maximum duration.

U-APSD

Unscheduled Automatic Power Save Delivery. U-APSD is a part of 802.11e and helps considerably in increasing the battery life of VoWLAN terminals.

UAM

Universal Access Method. UAM allows subscribers to access a wireless network after they successfully log in from a web browser.

UCC

Unified Communications and Collaboration. UCC is a term used to describe the integration of various communications methods with collaboration tools such as virtual whiteboards, real-time audio and video conferencing, and enhanced call control capabilities.

UDID

Unique Device Identifier. UDID is used to identify an iOS device.

UDP

User Datagram Protocol. UDP is a part of the TCP/IP family of protocols used for data transfer. UDP is typically used for streaming media. UDP is a stateless protocol, which means it does not acknowledge that the packets being sent have been received.

UDR

User Derivation Rule. UDR is a role assignment model used by the controllers running ArubaOS to assign roles and VLANs to the WLAN users based on MAC address, BSSID, DHCP-Option, encryption type, SSID, and the location of a user. For example, for an SSID with captive portal in the initial role, a UDR can be configured for scanners to provide a role based on their MAC OUI.

UHF

Ultra high frequency. UHF refers to radio frequencies between the range of 300 MHz and 3 GHz. UHF is also known as the decimeter band as the wavelengths range from one meter to one decimeter.

UI

User Interface.

UMTS

Universal Mobile Telecommunication System. UMTS is a third generation mobile cellular system for networks. See 3G.

UPnP

Universal Plug and Play. UPnp is a set of networking protocols that permits networked devices, such as personal computers, printers, Internet gateways, Wi-Fi APs, and mobile devices to seamlessly discover each other's presence on the network and establish functional network services for data sharing, communications, and entertainment.

URI

Uniform Resource Identifier. URI identifies the name and the location of a resource in a uniform format.

URL

Uniform Resource Locator. URL is a global address used for locating web resources on the Internet.

USB

Universal Serial Bus. USB is a connection standard that offers a common interface for communication between the external devices and a computer. USB is the most common port used in the client devices.

UTC

Coordinated Universal Time. UTC is the primary time standard by which the world regulates clocks and time.

UWB

Ultra-Wideband. UWB is a wireless technology for transmitting large amounts of digital data over a wide spectrum of frequency bands with very low power for a short distance.

VA

Virtual Appliance. VA is a pre-configured virtual machine image, ready to run on a hypervisor.

VBR

Virtual Beacon Report. VBR displays a report with the MAC address details and RSSI information of an AP.

VHT

Very High Throughput. IEEE 802.11ac is an emerging VHT WLAN standard that could achieve physical data rates of close to 7 Gbps for the 5 GHz band.

VIA

Virtual Intranet Access. VIA provides secure remote network connectivity for Android, Apple iOS, Mac OS X, and Windows mobile devices and laptops. It automatically scans and selects the best secure connection to the corporate network.

VLAN

Virtual Local Area Network. In computer networking, a single Layer 2 network may be partitioned to create multiple distinct broadcast domains, which are mutually isolated so that packets can only pass between them through one or more routers; such a domain is referred to as a Virtual Local Area Network, Virtual LAN, or VLAN.

VM

Virtual Machine. A VM is an emulation of a computer system. VMs are based on computer architectures and provide functionality of a physical computer.

VoIP

Voice over IP. VoIP allows transmission of voice and multimedia content over an IP network.

VoWLAN

Voice over WLAN. VoWLAN is a method of routing telephone calls for mobile users over the Internet using the technology specified in IEEE 802.11b. Routing mobile calls over the Internet makes them free, or at least much less expensive than they would be otherwise.

VPN

Virtual Private Network. VPN enables secure access to a corporate network when located remotely. It enables a computer to send and receive data across shared or public networks as if it were directly connected to the private network, while benefiting from the functionality, security, and management policies of the private network. This is done by establishing a virtual point-to-point connection through the use of dedicated connections, encryption, or a combination of the two.

VRD

Validated Reference Design. VRDs are guides that capture the best practices for a particular technology in field.

VRF

VisualRF. VRF is an AirWave Management Platform (AMP) module that provides a real-time, network-wide views of your entire Radio Frequency environment along with floor plan editing capabilities. VRF also includes overlays on client health to help diagnose issues related to clients, floor plan, or a specific location.

VRF Plan

VisualRF Plan. A stand-alone Windows client used for basic planning procedures such as adding a floor plan, provisioning APs, and generating a Bill of Materials report.

VRRP

Virtual Router Redundancy Protocol. VRRP is an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN.

VSA

Vendor-Specific Attribute. VSA is a method for communicating vendor-specific information between NASs and RADIUS servers.

VTP

VLAN Trunking Protocol. VTP is a Cisco proprietary protocol for propagating VLANs on a LAN.

W-CDMA

Wideband Code-Division Multiple Access. W-CDMA is a third-generation (3G) mobile wireless technology that promises much higher data speeds to mobile and portable wireless devices.

walled garden

Walled garden is a feature that allows blocking of unauthorized users from accessing network resources.

WAN

Wide Area Network. WAN is a telecommunications network or computer network that extends over a large geographical distance.

WASP

Wireless Application Service Provider. WASP provides a web-based access to applications and services that would otherwise have to be stored locally and makes it possible for customers to access the service from a variety of wireless devices, such as a smartphone or Personal Digital Assistant (PDA).

WAX

Wireless abstract XML. WAX is an abstract markup language and a set of tools that is designed to help wireless application development as well as portability. Its tags perform at a higher level of abstraction than that of other wireless markup languages such as HTML, HDML, WML, XSL, and more.

web service

Web services allow businesses to share and process data programmatically. Developers who want to provide integrated applications can use the API to programmatically perform actions that would otherwise require manual operation of the user interface.

WEP

Wired Equivalent Privacy. WEP is a security protocol that is specified in 802.11b and is designed to provide a WLAN with a level of security and privacy comparable to what is usually expected of a wired LAN.

WFA

Wi-Fi Alliance. WFA is a non-profit organization that promotes Wi-Fi technology and certifies Wi-Fi products if they conform to certain standards of interoperability.

Wi-Fi

Wi-Fi is a technology that allows electronic devices to connect to a WLAN network, mainly using the 2.4 GHz and 5 GHz radio bands. Wi-Fi can apply to products that use any 802.11 standard.

WIDS

Wireless Intrusion Detection System. WIDS is an application that detects the attacks on a wireless network or wireless system.

WiMAX

Worldwide Interoperability for Microwave Access. WiMAX refers to the implementation of IEEE 802.16 family of wireless networks standards set by the WiMAX forum.

WIP

Wireless Intrusion Protection. The WIP module provides wired and wireless AP detection, classification, and containment. It detects Denial of Service (DoS) and impersonation attacks, and prevents client and network intrusions.

WIPS

Wireless Intrusion Prevention System. WIPS is a dedicated security device or integrated software application that monitors the radio spectrum of WLAN network for rogue APs and other wireless threats.

WISP

Wireless Internet Service Provider. WISP allows subscribers to connect to a server at designated hotspots using a wireless connection such as Wi-Fi. This type of ISP offers broadband service and allows subscriber computers called stations, to access the Internet and the web from anywhere within the zone of coverage provided by the server antenna, usually a region with a radius of several kilometers.

WISPr

Wireless Internet Service Provider Roaming. The WISPr framework enables the client devices to roam between the wireless hotspots using different ISPs.

WLAN

Wireless Local Area Network. WLAN is a 802.11 standards-based LAN that the users access through a wireless connection.

WME

Wireless Multimedia Extension. WME is a Wi-Fi Alliance interoperability certification, based on the IEEE 802.11e standard. It provides basic QoS features to IEEE 802.11 networks. WMM prioritizes traffic according to four ACs: voice (AC_VO), video (AC_VI), best effort (AC_BE) and background (AC_BK). See WMM.

WMI

Windows Management Instrumentation. WMI consists of a set of extensions to the Windows Driver Model that provides an operating system interface through which instrumented components provide information and notification.

WMM

Wi-Fi Multimedia. WMM is also known as WME. It refers to a Wi-Fi Alliance interoperability certification, based on the IEEE 802.11e standard. It provides basic QoS features to IEEE 802.11 networks. WMM prioritizes traffic according to four ACs: voice (AC_VO), video (AC_VI), best effort (AC_BE), and background (AC_BK).

WPA

Wi-Fi Protected Access. WPA is an interoperable wireless security specification subset of the IEEE 802.11 standard. This standard provides authentication capabilities and uses TKIP for data encryption.

WPA2

Wi-Fi Protected Access 2. WPA2 is a certification program maintained by IEEE that oversees standards for security over wireless networks. WPA2 supports IEEE 802.1X/EAP authentication or PSK technology, but includes advanced encryption mechanism using CCMP that is referred to as AES.

WSDL

Web Service Description Language. WSDL is an XML-based interface definition language used to describe the functionality provided by a web service.

WSP

Wireless Service Provider. The service provider company that offers transmission services to users of wireless devices through Radio Frequency (RF) signals rather than through end-to-end wire communication.

WWW

World Wide Web.

X.509

X.509 is a standard for a public key infrastructure for managing digital certificates and public-key encryption. It is an essential part of the Transport Layer Security protocol used to secure web and email communication.

XAuth

Extended Authentication. XAuth provides a mechanism for requesting individual authentication information from the user, and a local user database or an external authentication server. It provides a method for storing the authentication information centrally in the local network.

XML

Extensible Markup Language. XML is a markup language that defines a set of rules for encoding documents in a format that is both human-readable and machine-readable.

XML-RPC

XML Remote Procedure Call. XML-RPC is a protocol that uses XML to encode its calls and HTTP as a transport mechanism. Developers who want to provide integrated applications can use the API to programmatically perform actions that would otherwise require manual operation of the user interface.

ZTP

Zero Touch Provisioning. ZTP is a device provisioning mechanism that allows automatic and quick provisioning of devices with a minimal or at times no manual intervention.