

53-1002487-01
November 2011



Brocade Mobility RFS4000, RFS6000, and RFS7000

System Reference Guide

Supporting software release 5.2.0.0 and later

BROCADE

Copyright © 2011 Brocade Communications Systems, Inc. All Rights Reserved.

Brocade, the B-wing symbol, BigIron, DCX, Fabric OS, FastIron, NetIron, SAN Health, ServerIron, and Turbolron are registered trademarks, and Brocade Assurance, Brocade NET Health, Brocade One, CloudPlex, MLX, VCS, VDX, and When the Mission Is Critical, the Network Is Brocade are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned are or may be trademarks or service marks of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

The authors and Brocade Communications Systems, Inc. shall have no liability or responsibility to any person or entity with respect to any loss, cost, liability, or damages arising from the information contained in this book or the computer programs that accompany it.

The product described by this document may contain "open source" software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <http://www.brocade.com/support/oscd>.

Brocade Communications Systems, Incorporated

Corporate and Latin American Headquarters
Brocade Communications Systems, Inc.
130 Holger Way
San Jose, CA 95134
Tel: 1-408-333-8000
Fax: 1-408-333-8101
E-mail: info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems China HK, Ltd.
No. 1 Guanghua Road
Chao Yang District
Units 2718 and 2818
Beijing 100020, China
Tel: +8610 6588 8888
Fax: +8610 6588 9999
E-mail: china-info@brocade.com

European Headquarters
Brocade Communications Switzerland Sàrl
Centre Swissair
Tour B - 4ème étage
29, Route de l'Aéroport
Case Postale 105
CH-1215 Genève 15
Switzerland
Tel: +41 22 799 5640
Fax: +41 22 799 5641
E-mail: emea-info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems Co., Ltd. (Shenzhen WFOE)
Citic Plaza
No. 233 Tian He Road North
Unit 1308 - 13th Floor
Guangzhou, China
Tel: +8620 3891 2000
Fax: +8620 3891 2111
E-mail: china-info@brocade.com

Document History

Title	Publication number	Summary of changes	Date
<i>Brocade Mobility RFS4000, RFS6000, and RFS7000 System Reference Guide</i>	53-1002312-01	New Additions for software version 5.1.0.0	June 2011
<i>Brocade Mobility RFS4000, RFS6000, and RFS7000 System Reference Guide</i>	53-1002487-01	New Additions for software version 5.2.0.0	November 2011

Contents

Chapter	About This Guide	
	In this chapter	xi
	Supported hardware and software	xi
	Document Conventions	xi
	Text formatting	xi
	Notes.	xii
	Related publications	xii
	Getting technical help.	xii
Chapter 1	Overview	
Chapter 2	Web UI Features	
	In this chapter	3
	Accessing the Web UI	3
	Browser and System Requirements	3
	Connecting to the Web UI	3
	Glossary of Icons Used	4
	Global Icons	5
	Dialog Box Icons.	6
	Table Icons	6
	Status Icons	7
	Configurable Objects	7
	Configuration Objects	10
	Configuration Operation Icons	10
	Access Type Icons	11
	Administrative Role Icons	11
	Device Icons	12
Chapter 3	Quick Start	
	In this chapter	13
	Using the Initial Setup Wizard	13
	Creating a managed WLAN.	22
	Assumptions.	22
	Design.	23
	Using the Controller GUI to Configure the WLAN	23
Chapter 4	Dashboard	
	In this chapter	79

Summary	79
Device Listing	80
System Screen	81
Network View	88

Chapter 5

Device Configuration

In this chapter	91
Basic Configuration	92
Basic Device Configuration	94
License Configuration	96
Assigning Certificates	98
Certificate Management	100
RSA Key Management	107
Certificate Creation	111
Generating a Certificate Signing Request	113
RF Domain Overrides	115
Profile Overrides	121
Controller Cluster Configuration Overrides (Controllers Only)	123
Access Point Adoption Overrides (Access Points Only)	125
Access Point Radio Power Overrides (Access Points Only)	127
Profile Interface Override Configuration	129
Overriding a Profile's Network Configuration	158
Overriding a Profile's Security Configuration	178
Overriding a Profile's Services Configuration	202
Overriding a Profile's Management Configuration	205
Overriding a Profile's Advanced Configuration	210
Auto Provisioning Policies	216
Configuring an Auto Provisioning Policy	218
Critical Resource Policy	222
Managing Critical Resource Policies	223
Managing Event Policies	224
Managing MINT Policies	225

Chapter 6

Wireless Configuration

In this chapter	227
Wireless LAN Policy	228
Basic WLAN Configuration	229
Configuring WLAN Security	232
WPA-TKIP Deployment Considerations	242
Configuring WLAN Firewall Support	251
Configuring Client Settings	256
Configuring WLAN Accounting Settings	258
Configuring Client Load Balancing Settings	260
Configuring Advanced WLAN Settings	261

Configuring WLAN QoS Policies	264
Configuring a WLAN's QoS WMM Settings	267
Configuring Rate Limit Settings	271
Configuring Multimedia Optimizations	276
Radio QoS Policy	278
Configuring Radio QoS Policies	280
Radio QoS Configuration and Deployment Considerations	288
AAA Policy	289
Association ACL	298
Association ACL Deployment Considerations	301
Smart RF Policy	301
Smart RF Configuration and Deployment Considerations	312

Chapter 7 Profile Configuration

In this chapter	313
General Profile Configuration	316
General Profile Configuration and Deployment Considerations	318
Profile Cluster Configuration (Controllers Only)	318
Controller Cluster Profile Configuration and Deployment Considerations	321
Profile Adoption Configuration (APs Only)	321
Profile Interface Configuration	322
Ethernet Port Configuration	323
Virtual Interface Configuration	330
Port Channel Configuration	334
Access Point Radio Configuration	340
WAN Backhaul Override Configuration	348
Profile Interface Deployment Considerations	350
Profile Network Configuration	350
Setting a Profile's DNS Configuration	350
ARP	352
Quality of Service (QoS)	353
Spanning Tree	355
Static Routes	358
Forwarding Database	359
Bridge VLAN	360
Cisco Discovery Protocol Configuration	364
Link Layer Discovery Protocol Configuration	365
Miscellaneous Network Configuration	367
Profile Network Configuration and Deployment Considerations	368

	Profile Security Configuration	368
	Defining Security Settings	369
	Setting the Certificate Revocation List (CRL) Configuration	370
	Configuring ISAKMP Policies	371
	Configuring Transform Sets	375
	Setting the Profile's VPN Configuration	377
	Setting the Profile's NAT Configuration	383
	Profile Security Configuration and Deployment Considerations	390
	Profile Services Configuration	391
	Profile Services Configuration and Deployment Considerations	393
	Profile Management Configuration	393
	Profile Management Configuration and Deployment Considerations	398
	Advanced Profile Configuration	399
	Configuring MINT	399
	Advanced Profile Miscellaneous Configuration	404
Chapter 8	RF Domain Configuration	
	In this chapter	407
	About RF Domains	407
	Default RF Domains	407
	User Defined RF Domains	408
	Managing RF Domains	408
	RF Domain Basic Configuration	410
	RF Domain Sensor Configuration	413
	RF Domain Overrides	414
	RF Domain Deployment Considerations	417
Chapter 9	Security Configuration	
	In this chapter	419
	Wireless Firewall	419
	Configuring a Firewall Policy	420
	Configuring IP Firewall Rules	431
	Configuring MAC Firewall Rules	434
	Firewall Deployment Considerations	437
	Wireless Client Roles	438
	Configuring a Client's Role Policy	438
	Intrusion Prevention	446
	Configuring a WIPS Policy	447
	Configuring an Advanced WIPS Policy	456
	Configuring a WIPS Device Categorization Policy	460
	Intrusion Detection Deployment Considerations	463
Chapter 10	Services Configuration	
	In this chapter	465

	Configuring Captive Portal Policies	465
	Configuring a Captive Portal Policy	465
	Creating DNS Whitelists	474
	Captive Portal Deployment Considerations	476
	Setting the Controller's DHCP Configuration	477
	Defining DHCP Pools	479
	Defining DHCP Server Global Settings	487
	DHCP Class Policy Configuration	489
	DHCP Deployment Considerations	490
	Setting the Controller's RADIUS Configuration	491
	Creating RADIUS Groups	492
	Defining User Pools	495
	Configuring RADIUS Server Policies	498
	RADIUS Deployment Considerations	509
Chapter 11	Management Access	
	In this chapter	511
	Viewing Management Access Policies	511
	Adding or Editing a Management Access Policy	514
	Management Access Deployment Considerations	525
Chapter 12	Diagnostics	
	In this chapter	527
	Fault Management	527
	Snapshots	532
	Core Snapshots	532
	Panic Snapshots	533
	Crash Files	534
	Advanced Diagnostics	536
	UI Debugging	536
Chapter 13	Operations	
	In this chapter	541
	Device Operations	541
	Managing Firmware and Config Files	541
	Managing File Transfers	544
	Using the Controller File Browser	546
	Using the AP Upgrade Browser	547
	Certificates	551
	Certificate Management	551
	RSA Key Management	559
	Certificate Creation	563
	Generating a Certificate Signing Request	564
	Smart RF	566
	Managing Smart RF for an RF Domain	567

Chapter 14

Statistics

In this chapter	571
System Statistics	571
Health	571
Inventory	573
Adopted Devices	575
Pending Adoptions	576
Licenses	577
RF Domain Statistics	579
.....	581
Access Point Statistics	593
Health	594
Inventory	596
Device	598
AP Upgrade	599
AP Detection	600
Wireless Client	601
Wireless LANs	603
Radios	604
Mesh	615
Interfaces	616
Network	621
Firewall	629
Certificates	638
WIPS	640
Sensor Servers	642
Captive Portal	643
Network Time	645
Load Balancing	648

Wireless Controller Statistics	649
Device Health	649
Inventory	651
Device	653
Cluster Peers	655
Adopted AP Statistics	656
AP Adoption History	657
Pending Adoptions	658
AP Detection	659
Wireless Clients	660
Wireless LANs	661
Critical Resource	663
Radios	664
Interfaces	668
Network	674
DHCP Server	682
Firewall	686
IPsec	694
Viewing Certificate Statistics	697
Controller WIPS Statistics	700
Advanced WIPS	702
Sensor Server	707
Captive Portal Statistics	708
Network Time	709
Wireless Client Statistics	713
Health	713
Details	716
Traffic	718
WMM TSPEC	720

Chapter A

Publicly Available Software

In this appendix	723
General Information	723
Open Source Software Used	724
Wireless Controller	724
AP650	725
AP51xx	727
AP7131	727
OSS Licenses	729
GNU General Public License 2.0	729
GNU Lesser General Public License 2.1	734
BSD Style Licenses	741
MIT License	742
WU-FTPD License	742
Open SSL License	743
ZLIB License	745
Open LDAP Public License	745
Apache License 2.0	746
Drop Bear License	749
Sun Community Source License	750

Chapter B

Upgrading from Version 4 software to Version 5

- In this appendix 763
- AP650 763
 - Upgrade Version 4.x to 5.x 763
 - Downgrade Version 5.x to 4.x 763
 - AP650 Connectivity in Version 5.x 763
- AP7131 / AP71xx. 764
 - Upgrade version 4.x to 5.x. 764
 - Downgrade 5.x to 4.x. 764
 - Migration Files 764
 - Recovery of a Pre-4.1.1-18R Access Point. 765
 - AP7131 / 7131N Connectivity in version 5.x 765
 - Configuration Restoration 765
 - Version 5.1 Minimum Requirements 765
- Controllers 766
 - Upgrade Version 4.x to 5.x 766
 - Migrating Version 4.x Configuration File to 5.2 Configuration File 766
 - Downgrade Version 5.x to 4.x 766
 - Version 5.2 Minimum Requirements 766

About This Guide

In this chapter

- [Supported hardware and software](#) xi
- [Document Conventions](#) xi
- [Related publications](#) xii
- [Getting technical help](#) xii

Supported hardware and software

This guide provides information on using the following Brocade wireless controllers and access points:

- Brocade Mobility RFS7000 Controller
- Brocade Mobility RFS6000 Controller
- Brocade Mobility RFS4000 Controller
- Brocade Mobility 7131 Series Access Point
- Brocade Mobility 650 Access Point
- Brocade Mobility 6511 Access Point

Document Conventions

This section describes text formatting conventions and important notice formats used in this document.

Text formatting

The narrative-text formatting conventions that are used are as follows:

bold text	Identifies command names
	Identifies the names of user-manipulated GUI elements
	Identifies keywords
	Identifies text to enter at the GUI or CLI
<i>italic text</i>	Provides emphasis
	Identifies variables
	Identifies document titles
code text	Identifies CLI output

For readability, command names in the narrative portions of this guide are presented in bold; for example, **show version**.

Notes

The following notice statement is used in this manual.

NOTE

A note provides a tip, guidance or advice, emphasizes important information, or provides a reference to related information.

Related publications

The following Brocade Communications Systems, Inc. documents supplement the information in this guide and can be located at <http://www.brocade.com/ethernetproducts>.

- Installation Guides - Each controller has a unique installation guide which describes the basic hardware setup and configuration required to transition to more advanced configuration of the wireless controllers.
- *Brocade Mobility RFS4000, RFS6000, and RFS7000 System Reference Guide* (this document) - Describes configuration of the Brocade wireless controllers using the Web UI.
- *Brocade Mobility RFS4000, RFS6000 and RFS7000 CLI Reference Guide* - Describes the Command Line Interface (CLI) and Management Information Base (MIB) commands used to configure the Brocade controllers.

If you find errors in the guide, send an e-mail to documentation@brocade.com.

Getting technical help

To contact Technical Support, go to <http://www.brocade.com/services-support/index.page> for the latest e-mail and telephone contact information.

Overview

The Brocade Mobility family of wireless controllers with the 802.11n Access Points enable the centralized distribution of high performance, secure and resilient wireless voice and data services to remote locations with the scalability required to meet the needs of large distributed enterprises.

A Brocade Mobility controller provides a single platform capable of delivering wireless voice and data inside and outside the enterprise for small, medium and large enterprise deployments. Improve operational efficiency and reduce the cost of mobility with a powerful comprehensive feature set including adaptive AP, which delivers unmatched performance, security, reliability and scalability to enable networks for business mobility at a low TCO.

Wireless controllers provide local centralized management and control of 802.11n Access Points and the Brocade Mobility provide the necessary core switching and routing to eliminate additional routing and switching infrastructure.

802.11n is the next generation WLAN standard that provides improved performance and coverage compared with previous 802.11 specifications. 802.11n provides enhancements to support throughput up to 450 Mbps. With these enhancements Brocade Mobility next generation 802.11n Access Points offer client data-rates of up to 300Mbps.

The Brocade Mobility solution uses 802.11n Access Points and peer controllers to adapt to the dynamic circumstances of their deployment environment. The Brocade Mobility architecture provides a customized site-specific deployment, supporting the best path and routes based on the user, location, the application and the best route available (both wireless and wired). The Brocade Mobility solution assures end-to-end quality, reliability and security without latency and performance degradation. The Brocade Mobility supports rapid application delivery, mixed-media application optimization and quality assurance.

The Brocade Mobility architecture is designed for 802.11n networking. It leverages the best aspects of independent and dependent architectures to create a smart network that meets the connectivity, quality and security needs of each user deployment and their application requirements, based on the availability of network resources, including wired networks.

By distributing intelligence and control between the wireless controllers and Access Points, Brocade Mobility can route data directly using the best path, as determined by factors including the user, the location, the application and available wireless and wired resources. As a result, the additional load placed on the wired network from 802.11n is significantly reduced, as traffic does not require an unnecessary backhaul to a central controller.

Within a Brocade Mobility, up to 80% of the network traffic can remain on the wireless mesh, and never touch the wired network, so the 802.11n load impact on the wired network is negligible. In addition, latency and associated costs are reduced while reliability and scalability are increased. Brocade Mobility enables the creation of dynamic wireless traffic flows, so any bottleneck is avoided, and the destination is reached without latency or performance degradation. This behavior delivers a significantly better quality of experience for the end user.

The same distributed intelligence enables more resilience and survivability, since the Access Points keep users connected and traffic flowing with full QoS, security and mobility even if the connection to the wireless controller is interrupted due to a wired network or backhaul problem.

Even when the network is fully operational, outside RF interference sources or unbalanced wireless network loading can be automatically corrected by the Brocade Mobility Smart RF system. Smart RF senses interference or potential client connectivity problems and makes the required changes to channel and Access Point radio power while minimizing the impact to latency sensitive applications like VoIP. Using Smart RF, the managed network can continuously adjust Access Point power and channel assignments for self-recovery if an AP fails or a coverage hole is detected.

Additionally, integrated Access Point sensors in conjunction with AirDefense Network Assurance alerts administrators of interference and network coverage problems, which shortens response times and boosts overall reliability and availability of the Brocade Mobility.

Network traffic optimization protects Brocade Mobility from broadcast storms and minimizes congestion on the wired network. Brocade Mobility solutions provide VLAN load balancing, WAN traffic shaping and optimizations in *dynamic host configuration protocol* (DHCP) responses and *Internet group management protocol* (IGMP) snooping for multicast traffic flows in wired and wireless networks. Thus, users benefit from an extremely reliable network that adapts to meet their needs and delivers mixed-media applications.

Firmware and configuration updates are supported within the managed network, from one Access Point to another, over the air or wire, and can be centrally managed by the controller. Controllers no longer need to push firmware and configurations to each individual Access Point, reducing unnecessary network congestion.

Brocade Mobility uses remote Authentication Dial-in User Service (RADIUS) synchronization capabilities between the core and the access layer. If the central authentication mechanism is not available, users can authenticate using the controller local RADIUS resources, and continue network support with secure access.

Web UI Features

In this chapter

- [Accessing the Web UI](#) 3
- [Glossary of Icons Used](#) 4

Brocade Mobility contains a Web UI allowing network administrators to manage and view Brocade Mobility Wireless controller settings, configuration and status. This *Graphical User Interface* (GUI) allows full control of all managed features.

Wireless controllers also include a *Command Line Interface* (CLI) for managing and viewing settings, configuration and status. For more information on the command line interface and a full list of available commands, see the *Brocade Mobility RFS4000, RFS6000, and RFS7000 CLI Reference Guide* available at <http://www.brocade.com>

Accessing the Web UI

Brocade Mobility wireless controllers use a UI accessed using any supported Web browser on a client connected to the subnet the Web UI is configured on.

Browser and System Requirements

To access the UI, a browser supporting Flash Player 10 is recommended. The system accessing the GUI should have a minimum of 512Mb or RAM for the UI to display and function properly. The UI is based on Flex, and does not use Java as its underlying framework.

The following browsers have been validated with the Web UI:

- Firefox 3.6
- Internet Explorer 7.x
- Internet Explorer 8.x

NOTE

Throughout the Web UI leading and trailing spaces are not allowed in any text fields. In addition, the “?” character is also not supported in text fields.

Connecting to the Web UI

1. Connect one end of an Ethernet cable to any of the five LAN ports on the front of the RFS4011 and connect the other end to a computer with a working Web browser.
2. Set the computer to use an IP address between 192.168.0.10 and 192.168.0.250 on the connected port. Set a subnet / network mask of 255.255.255.0.

3. Once the computer has an IP address, point the Web browser to: <https://192.168.0.1/> and the following login screen will display.

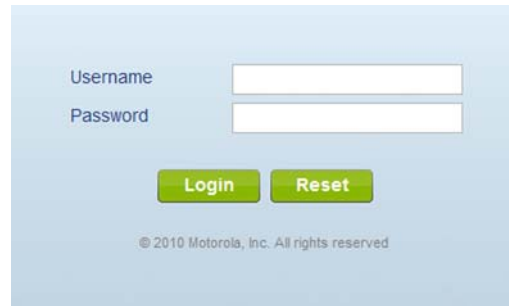


FIGURE 1 Web UI Login Screen

4. Enter the default username *admin* in the **Username** field.
5. Enter the default password *admin123* in the **Password** field.
6. Click the **Login** button to load the management interface.
7. If this is the first time the UI has been accessed, a dialogue displays to begin an initial setup wizard. For more information on using the initial setup wizard see [Using the Initial Setup Wizard on page 3-13](#).

Glossary of Icons Used








The UI uses a number of icons used to interact with the system, gather information, and obtain status for the entities managed by the system. This chapter is a compendium of the icons used. This chapter is organized as follows:

- [Global Icons](#)
- [Dialog Box Icons](#)
- [Table Icons](#)
- [Status Icons](#)
- [Configurable Objects](#)
- [Configuration Objects](#)
- [Configuration Operation Icons](#)
- [Access Type Icons](#)
- [Administrative Role Icons](#)
- [Device Icons](#)

Global Icons

Glossary of Icons Used

This section lists global icons available throughout the controller interface.

	<i>Logoff</i> – Select this icon to log out of the managed system. This icon is always available and is located at the top right corner of the UI.
	<i>Add</i> – Select this icon to add a row in a table. When selected, a new row is created in the table or a dialog box displays where you can enter values for a particular list.
	<i>Delete</i> – Select this icon to remove a row from a table. When selected, the selected row is deleted.
	<i>More Information</i> – Select this icon to display a pop up with supplementary information that may be available for an item.
	<i>Trash</i> – Select this icon to remove a row from a table. When selected, the row is immediately deleted.
	<i>Create new policy</i> – Select this icon to create a new policy. Policies define different configuration parameters that can be applied to individual device configurations, device profiles and RF Domains.
	<i>Edit policy</i> – Select this icon to edit an existing policy. To edit a policy, select a policy and this icon.

Dialog Box Icons

Glossary of Icons Used

These icons indicate the current state of various controls in a dialog. These icons enables you to gather, at a glance, the status of all the controls in a dialog. The absence of any of these icons next to a control indicates the value in that control has not been modified from its last saved configuration.







	<i>Entry Updated</i> – Indicates a value has been modified from its last saved configuration.
	<i>Entry Update</i> – States that an override has been applied to a device's profile configuration.
	<i>Mandatory Field</i> – Indicates this control value is a mandatory configuration item. You will not be allowed to proceed further without providing all mandatory values in this dialog.
	<i>Error in Entry</i> – Indicates there is an error in a value entered in this control. A small red popup provides a likely cause of the error.

Table Icons

Glossary of Icons Used






The following two override icons are status indicators for transactions:

	<i>Table Row Overridden</i> – Indicates a change (profile configuration override) has been made to a table row and the change will not be implemented until saved. This icon represents a change from this device's profile assigned configuration.
	<i>Table Row Added</i> – Indicates a new row has been added to a table and the change is not implemented until saved. This icon represents a change from this device's profile assigned configuration.

Status Icons

[Glossary of Icons Used](#)





These icons indicate device status, operations on the wireless controller, or any other action that requires a status returned to the user.













	<i>Fatal Error</i> – States there is an error causing a managed device to stop functioning.
	<i>Error</i> – Indicates an error exists requiring intervention. A managed action has failed, but the error is not system wide.
	<i>Warning</i> – States a particular action has completed, but errors were detected that did not prevent the process from completing. Intervention might still be required to resolve subsequent warnings.
	<i>Success</i> – Indicates everything is well within the managed network or a process has completed successfully without error.
	<i>Information</i> – This icon always precedes information displayed to the user. This may either be a message displaying progress for a particular process, or just be a message from the system.











Configurable Objects

[Glossary of Icons Used](#)

These icons represent configurable items within the controller's UI.

	<i>Device Configuration</i> – Represents a configuration file supporting a device category (AP, Wireless Controller etc.).
	<i>Provisioning Policy</i> – Represents a provisioning policy. Adoption policies are a set of configuration parameters that define how APs and wireless clients are adopted by a controller.
	<i>Critical Resource Policy</i> – States a critical resource policy has been applied. Critical resources are resources whose availability is essential to the managed network. If any of these resources is unavailable, an administrator is notified.
	<i>Wireless LANs</i> – States an action impacting a managed WLAN has occurred.







	<i>WLAN QoS Policy</i> – States a <i>quality of service policy</i> (QoS) configuration has been impacted.
	<i>Radio QoS Policy</i> – Indicates a radio's QoS configuration has been impacted.
	<i>AAA Policy</i> – Indicates an <i>Authentication, Authorization and Accounting</i> (AAA) policy has been impacted. AAA policies define RADIUS authentication and accounting parameters.
	<i>Association ACL</i> – Indicates an <i>Access Control List</i> (ACL) configuration has been impacted. An ACL is a set of configuration parameters either allowing or denying access to managed resources.
	<i>Smart RF Policy</i> – States a Smart RF policy has been impacted. Smart RF enables neighboring APs to take over for an AP if it becomes unavailable. This is accomplished by increasing the power of radios on nearby APs to cover the hole created by the non-functioning AP.
	<i>Profile</i> – States a device profile configuration has been impacted. A profile is a collection of configuration parameters used to configure a device or a feature.
	<i>Bridging Policy</i> – Indicates a bridging policy configuration has been impacted. A Bridging Policy defines which VLANs are bridged, and how local VLANs are bridged between the wired and wireless sides of the managed network.
	<i>RF Domain</i> – States an RF Domain configuration has been impacted.
	<i>Firewall Policy</i> – Indicates a firewall policy has been impacted. Firewalls provide a barrier that prevents unauthorized access to resources while allowing authorized access to external and internal controller resources.
	<i>IP Firewall Rules</i> – Indicates an IP firewall rule has been applied. An IP based firewall rule implements restrictions based on the IP address in a received packet.
	<i>MAC Firewall Rules</i> – States a MAC based firewall rule has been applied. A MAC based firewall rule implements firewall restrictions based on the MAC address in a received packet.
	<i>Wireless Client Role</i> – Indicates a wireless client role has been applied to a managed client. The role could be either sensor or client.

	<i>WIPS Policy</i> – States the conditions of a WIPS policy have been invoked. WIPS prevents unauthorized access to the managed network by checking for (and removing) rogue APs and wireless clients.
	Advanced WIPS Policy – States the conditions of an advanced WIPS policy have been invoked.
	<i>Device Categorization</i> – Indicates a device categorization policy has been applied. This is used by the intrusion prevention system to categorize APs or wireless clients as either sanctioned or unsanctioned devices. This enables devices to bypass the intrusion prevention system.
	<i>Captive Portal</i> – States a captive portal is being applied. Captive portal is used to provide hotspot services to wireless clients.
	<i>DNS Whitelist</i> – A DNS whitelist is used in conjunction with captive portal to provide hotspot services to wireless clients.
	<i>DHCP Server Policy</i> – Indicates a DHCP server policy is being applied. DHCP provides IP addresses to wireless clients. A DHCP server policy configures how DHCP provides IP addresses.
	<i>RADIUS Group</i> – Indicates the configuration of RADIUS group has been defined and applied. A RADIUS group is a collection of RADIUS users with the same set of permissions.
	<i>RADIUS User Pools</i> – States a RADIUS user pool has been applied. RADIUS user pools are a set of IP addresses that can be assigned to an authenticated RADIUS user.
	<i>RADIUS Server Policy</i> – Indicates a RADIUS server policy has been applied. A RADIUS server policy is a set of configuration attributes used when a RADIUS server is configured for AAA.
	<i>Management Policy</i> – Indicates a management policy has been applied. Management policies configure access control, authentication, traps and administrator permissions.

Configuration Objects

[Glossary of Icons Used](#)




These configuration icons are used to define the following:

	<i>Configuration</i> – Indicates an item capable of being configured by a controller interface.
	<i>View Events / Event History</i> – Defines a list of events. Click this icon to view events or view the event history.
	<i>Core Snapshots</i> – Indicates a core snapshot has been generated. A core snapshot is a file that records status when a process fails on the wireless controller.
	<i>Panic Snapshots</i> – Indicates a panic snapshot has been generated. A panic snapshot is a file that records status when the wireless controller fails without recovery.
	<i>UI Debugging</i> – Select this icon/link to view current NETCONF messages.
	<i>View UI Logs</i> – Select this icon/link to view the different logs generated by the UI, FLEX and the error logs.

Configuration Operation Icons

[Glossary of Icons Used](#)





The following operations icons are used to define configuration operations:

	<i>Revert</i> – When selected, any changes made after the last saved configuration are restored back to the last saved configuration.
	<i>Commit</i> – When selected, all changes made to the configuration are written to the system. Once committed, changes cannot be reverted.
	<i>Save</i> – When selected, changes are saved to the configuration.

Access Type Icons

[Glossary of Icons Used](#)




The following icons display a user access type:





	<i>Web UI</i> – Defines a Web UI controller access permission. A user with this permission is permitted to access an associated device's Web UI.
	<i>Telnet</i> – Defines a TELNET access permission. A user with this permission is permitted to access an associated device using TELNET.
	<i>SSH</i> – Indicates a SSH access permission. A user with this permission is permitted to access an associated device using SSH.
	<i>Console</i> – Indicates a console access permission. A user with this permission is permitted to access an associated device using the device's serial console.

Administrative Role Icons

[Glossary of Icons Used](#)

The following icons identify the different administrative roles allowed on the system:

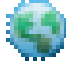





	<i>Superuser</i> – Indicates superuser privileges. A superuser has complete access to all configuration aspects of the connected device.
	<i>System</i> – States system user privileges. A system user is allowed to configure general settings, such as boot parameters, licenses, auto install, image upgrades etc.
	<i>Network</i> – Indicates network user privileges. A network user is allowed to configure wired and wireless parameters, such as IP configuration, VLANs, L2/L3 security, WLANs and radios.

	<i>Security</i> – Indicates security user privileges. A security level user is allowed to configure all security related parameters.
	<i>Monitor</i> – Defines a monitor role. This role provides no configuration privileges. A user with this role can view the system configuration but cannot modify it.
	<i>Help Desk</i> – Indicates help desk privileges. A help desk user is allowed to use troubleshooting tools like sniffers, execute service commands, view or retrieve logs and reboot the controller.
	<i>Web User</i> – Indicates a web user privilege. A Web user is allowed accessing the device's Web UI.

Device Icons

[Glossary of Icons Used](#)

The following icons represent the different device types managed by the system:

	<i>System</i> – This icon represents the entire Brocade Mobility supported system.
	<i>Cluster</i> – This icon represents a cluster. A cluster is a set of wireless controllers working collectively to provide redundancy and load sharing.
	<i>Wireless Controller</i> – This icon indicates a RFS6000 or a RFS7000 wireless controller that's part of the managed network.
	<i>Wireless Controller</i> – This icon indicates a RFS4000 wireless controller that's part of the managed network.
	<i>Access Point</i> – This icon indicates any access point that's part of the managed network.
	<i>Wireless Client</i> – This icon defines any wireless client connection within the managed network.

Quick Start

In this chapter

- [Using the Initial Setup Wizard](#)..... 13
- [Creating a managed WLAN](#)..... 22

Brocade Mobility supported controllers utilize an initial settings wizard to streamline the process of getting the controller on the network for the first time. The wizard defines configure location, network and WLAN settings and assists in discovery of access points. For instructions on how to use the initial setup wizard, see *Using the Initial Setup Wizard on page 3-13*.

Using the Initial Setup Wizard

Once the controller is deployed and powered on, complete the following to get up and running and access management functions:

Connect one end of an Ethernet cable to a port on the front of the controller and connect the other end to a computer with a working Web browser.

Set the computer to use an IP address between 192.168.0.10 and 192.168.0.250 on the connected port. Set a subnet/network mask of 255.255.255.0.

Once the computer has an IP address, point the Web browser to: <https://192.168.0.1/>, the following login screen will display.

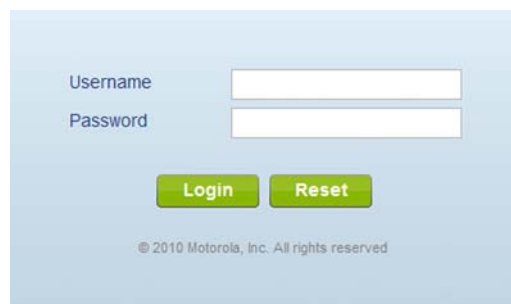


FIGURE 2 Web UI Login screen

Enter the default username *admin* in the **Username** field.

Enter the default password *admin123* in the **Password** field.

Select the **Login** button to load the management interface.

If this is the first time the management interface has been accessed, a dialogue displays to start the initial setup wizard. Select the **Start Wizard** button.



FIGURE 3 Initial Setup Wizard

Change the default **Password** and enter a **Location**, and **Contact** name. Select a **Time Zone** and **Country** for the controller.

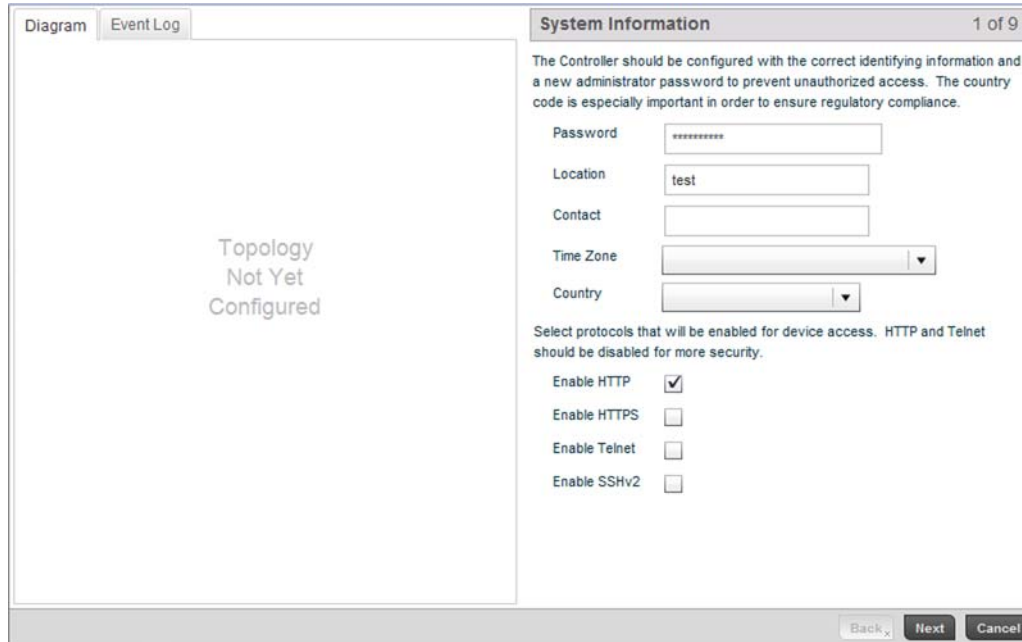


FIGURE 4 Initial Setup Wizard - System Information screen

Select each of the protocols (access methods) you would like to permit for the controller. Select the **Next** button to continue to the **Topology Selection** screen.

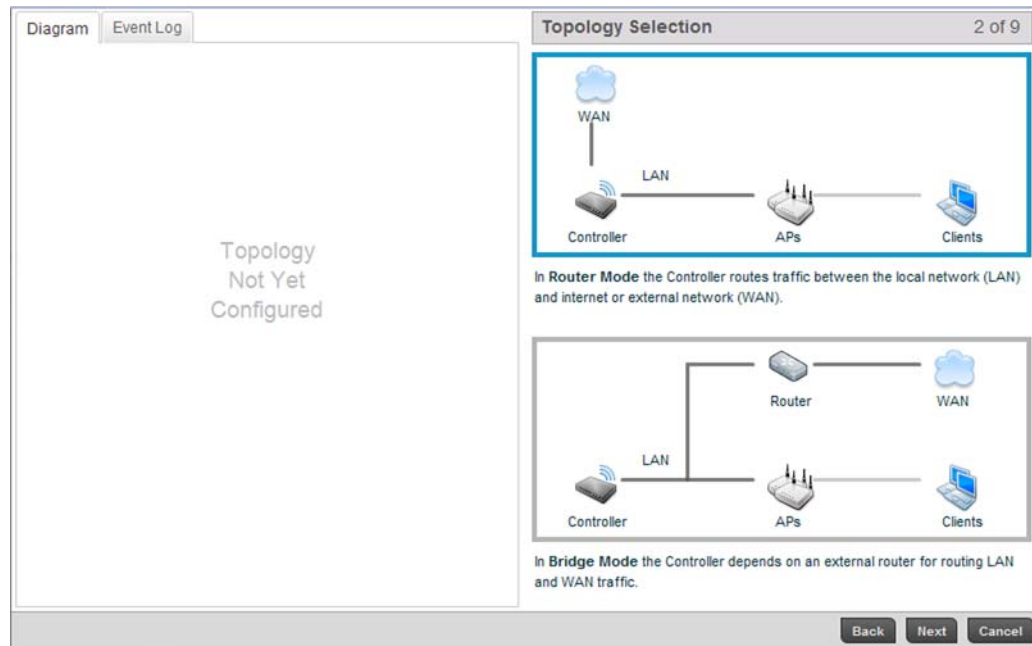


FIGURE 5 Initial Setup Wizard - Topology Selection screen

Select a network topology for the controller deployment. The mode selected will result in a specific screen flow for the remainder of the initial setup wizard.

Router Mode Using Router Mode, the controller routes traffic between the local network (LAN) and internet or external network (WAN).

Bridge Mode Using Bridge Mode, the controller uses an external router for LAN and WAN traffic. Routing is generally used for one device, whereas bridging is typically used with a larger density network.

For the purposes of this example, select **Router Mode**.

Select the **Next** button to continue to the **LAN Configuration** screen.

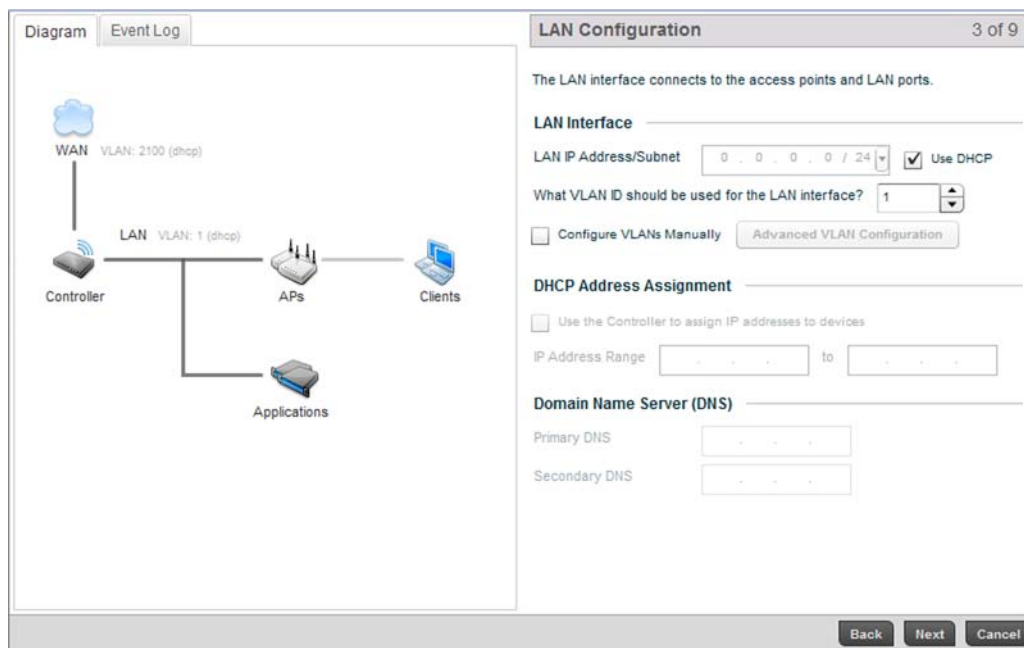


FIGURE 6 Initial Setup Wizard - LAN Configuration screen

The **LAN Configuration** screen is partitioned into **LAN Interface**, **DHCP Address Assignment** and **Domain Name Server (DNS)** fields.

- a. Refer to the **LAN Interface** field for LAN IP address, subnet and VLAN configuration parameters.

LAN IP Address / Subnet Enter an IP Address and a subnet for the controller's LAN interface. If the Use DHCP checkbox is enabled, this field will not be configurable.

Use DHCP Select the Use DHCP checkbox to enable automatic network configuration using a DHCP server. If this option is enabled, the LAN IP Address/Subnet, DHCP Address Assignment and Domain Name fields are populated by the DHCP server.

What VLAN ID should be used for the LAN interface Set the VLAN ID to associate with the LAN Interface. The default setting is VLAN 1.

Configure VLANs Manually Select the Configure VLANs Manually checkbox to enable advanced manual VLAN configuration.
For more information, see [Virtual Interface Configuration on page 7-330](#).

Advanced VLAN Configuration Select the Advanced VLAN Configuration button to set associations between VLANs and controller physical interfaces.

For the purposes of this example, select **Use DHCP** and uncheck **Configure VLANs Manually**.

- b. Refer to the **DHCP Address Assignment** field to set DHCP server settings for the LAN interface.

Use the Controller to assign IP addresses to devices

Select the **Use the Controller to assign IP addresses to devices** checkbox to enable the onboard DHCP server to provide IP and DNS information to clients on the LAN interface.

IP Address Range

Enter a starting and ending IP Address range for client assignments on the LAN interface. It's good practice to avoid assigning IP addresses from x.x.x.1 - x.x.x.10 and x.x.x.255 as they are often reserved for standard network services.

- c. Refer to the **Domain Name Server (DNS)** field to set DNS server settings on the LAN interface.

Primary DNS

Enter an IP Address for the main DNS server for the controller LAN interface.

Secondary DNS

Enter an IP Address for the backup DNS server for the LAN interface.

Select **Next** to save the LAN configuration settings and move to the WAN Configuration screen.

The WAN Configuration screen is partitioned into **WAN Interface** and **Gateway** fields.

The screenshot shows the WAN Configuration screen in the Initial Setup Wizard. On the left, a network diagram illustrates the Controller's connections to the WAN (VLAN: 2100, 157.235.8.11/24) and LAN (VLAN: 1, dhcp) interfaces, with the LAN serving APs, Clients, and Applications. On the right, the configuration fields are as follows:

- WAN Interface:** WAN IP Address/Subnet is 157.235.8.11/24. The 'Use DHCP' checkbox is unchecked. The VLAN ID is set to 2100, and the port connected to the external network is up1. The 'Enable NAT on the WAN Interface' checkbox is also unchecked.
- Gateway:** The Default Gateway is set to 157.235.8.1.
- Commit:** A warning icon and text state: "Clicking 'Next/Commit' will commit changes to the Controller. This may take several minutes and affect connectivity to the device, requiring the UI to be closed and re-launched. Connection or login issues may occur if network or access parameters have changed."

At the bottom of the screen, there are three buttons: Back, Next/Commit, and Cancel.

FIGURE 7 Initial Setup Wizard - WAN Configuration screen

- a. Refer to the **WAN Interface** field to set the WAN IP address, subnet and VLAN configuration.

WAN IP Address/Subnet	Enter an IP Address and a subnet for the controller's WAN interface. If the Use DHCP checkbox is enabled, this field will not be configurable.
Use DHCP	Select the Use DHCP checkbox to enable an automatic network configuration using a DHCP Server. If this option is enabled, the WAN IP Address/Subnet and Gateway fields are populated by the DHCP server.
What VLAN ID should be used for the WLAN interface	Set the VLAN ID to associate with the WLAN Interface. The default setting is VLAN 2100. For more information, see Virtual Interface Configuration on page 7-330 .
What port is connected to the external network?	Select the physical controller port connected to the WAN interface. The list of available ports varies based on controller model.
Enable NAT on the WAN Interface	Click the Enable NAT on WAN Interface checkbox to enable <i>Network Address Translation</i> (NAT) allowing traffic to pass between the controller WAN and LAN interfaces.

- b. Refer to the **Gateway** field to set the **Default Gateway**.

Default Gateway	Enter an IP Address for the controller's default gateway on the WAN interface. If the Use DHCP checkbox is enabled, this field will not be configurable.
------------------------	--

Select **Next** to save the WAN configuration settings and move to the WLAN Setup screen.

Use The WLAN Setup screen to enable managed WLANs.

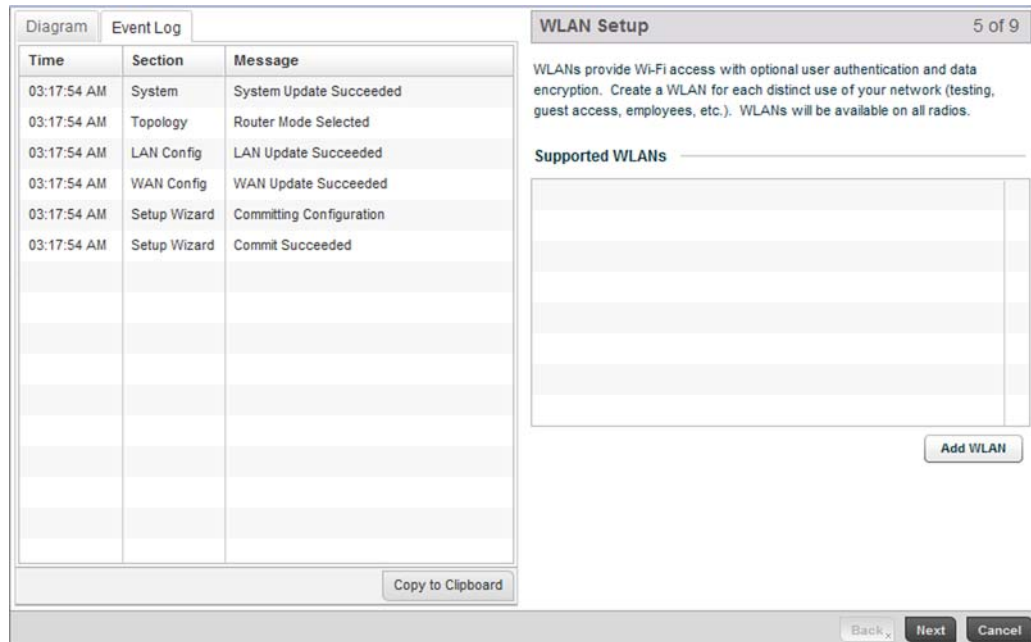


FIGURE 8 Initial Setup Wizard - WLAN Setup screen

Select the **Add WLAN** button to display a screen used to define WLAN settings.

FIGURE 9 Initial Setup Wizard

Set the following parameters for new WLAN configurations:

- | | |
|------------------|--|
| SSID | Enter or modify the <i>Services Set Identification</i> (SSID) associated with the WLAN. The WLAN name is auto-generated using the SSID until changed by the user. The maximum number of characters is 32. Do not use any of these characters (< > " & \ ? ,). |
| WLAN Type | Select a basic authentication and encryption scheme for the WLAN. Available options include:
No authentication, no encryption
Captive portal authentication, no encryption
PSK authentication, WPA2 encryption
EAP authentication, WPA2 encryption
For more information, see Configuring WLAN Security on page 6-232 . |
| VLAN Id | Select a VLAN to associate with WLAN traffic. Each configured VLAN is available for selection. |
| WPA Key | Enter either an alphanumeric string of 8 to 63 ASCII characters or 64 HEX characters as the primary string both transmitting and receiving authenticators must share. The alphanumeric string allows character spaces. The wireless controller converts the string to a numeric value. This passphrase saves the administrator from entering the 256-bit key each time keys are generated. |

Select **OK** to exit, then select **Next** to continue to the RADIUS Authentication screen.

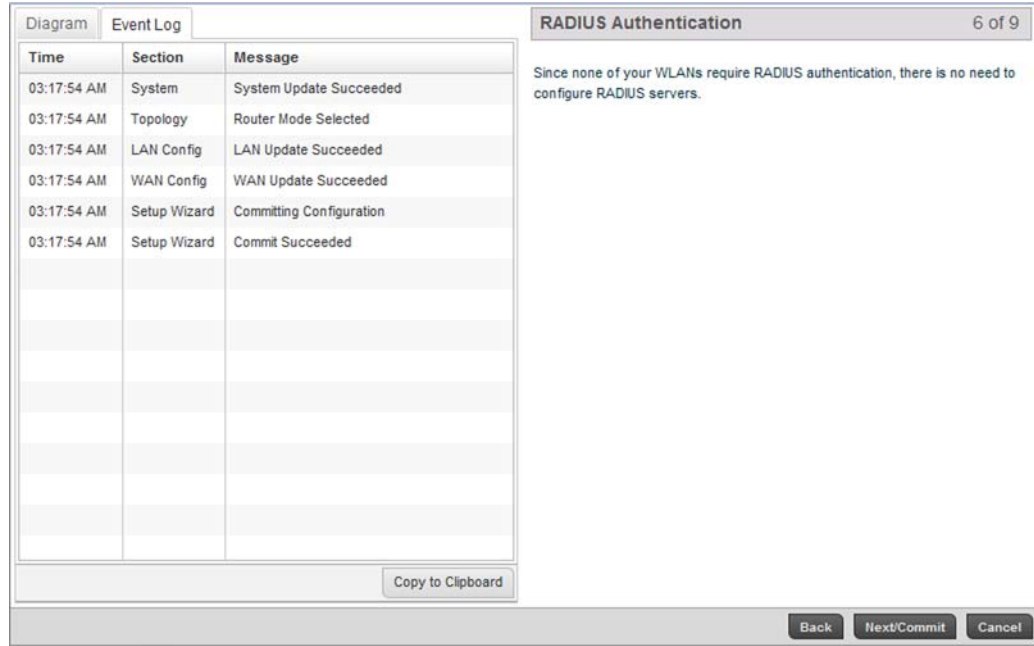


FIGURE 10 Initial Setup Wizard - RADIUS Authentication screen

Within this example, there's no action required since no WLANs require RADIUS authentication. Select **Next/Commit** to continue to the AP Discovery screen.

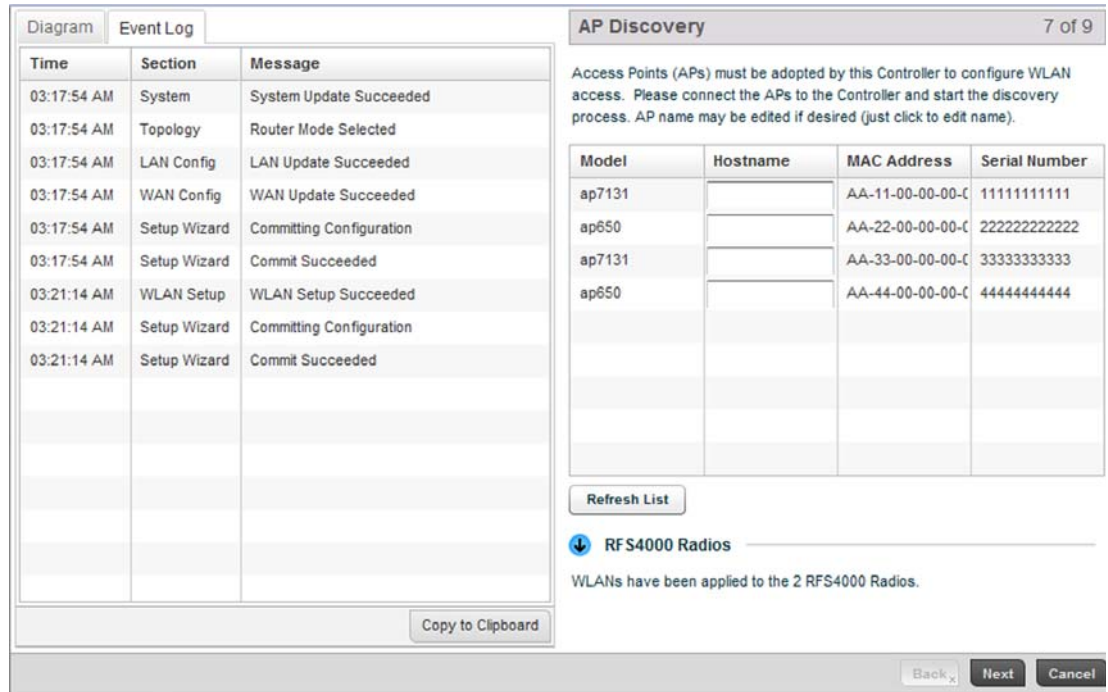


FIGURE 11 Initial Setup Wizard - AP Discovery screen

The AP Discovery screen displays a list of Access Points discovered by the controller. If you have connected any APs to the controller select the **Refresh List** button to update the list of known APs.

Optionally, define a **Hostname** for each known AP.

NOTE

If using a RFS4011 model controller, configured WLANs are automatically applied to the internal radios.

Select the **Next** button to move to the **Wireless Client Association** screen.

The screenshot shows the 'Wireless Client Association' screen (8 of 9) in the Initial Setup Wizard. On the left, there is an 'Event Log' table with columns for Time, Section, and Message. The main area contains a table of WLAN configurations and a list of associated wireless clients. Below the table are buttons for 'Refresh' and 'Save'. A note at the bottom states: 'Clicking Finish/Save will persistently store the configuration on the Controller. This means it will remain even if the Controller is restarted.' At the bottom right, there are buttons for 'Back', 'Finish/Save', and 'Cancel'.

Time	Section	Message
03:17:54 AM	System	System Update Succeeded
03:17:54 AM	Topology	Router Mode Selected
03:17:54 AM	LAN Config	LAN Update Succeeded
03:17:54 AM	WAN Config	WAN Update Succeeded
03:17:54 AM	Setup Wizard	Committing Configuration
03:17:54 AM	Setup Wizard	Commit Succeeded
03:21:14 AM	WLAN Setup	WLAN Setup Succeeded
03:21:14 AM	Setup Wizard	Committing Configuration
03:21:14 AM	Setup Wizard	Commit Succeeded

WLAN	Username	MAC Address	IP Address
wlan1	user1	00:21:6A:15:17:9E	10.1.1.1
wlan1	user1	00:21:6A:15:17:9E	10.1.1.1
wlan1	user1	AA-11-11-00-00-C	10.1.1.1
wlan2	user1	AA-11-22-00-00-C	10.1.1.2
wlan3	user1	AA-11-33-00-00-C	10.1.1.3
wlan4	user4	AA-11-44-00-00-C	10.1.1.4
wlan5	user5	AA-11-55-00-00-C	10.1.1.5
wlan6	user6	AA-11-66-00-00-C	10.1.1.6
wlan7	user7	AA-11-77-00-00-C	10.1.1.7

FIGURE 12 Initial Setup Wizard - Wireless Client Association screen

The **Wireless Client Association** screen displays adopted clients and the WLANs they are associated with.

To verify the WLAN configuration, associate a wireless client with each configured WLAN. After associating, click the **Refresh** button to update the list of associated wireless clients.

Select the **Finish/Save** button to complete the wizard and display a summary of changes.

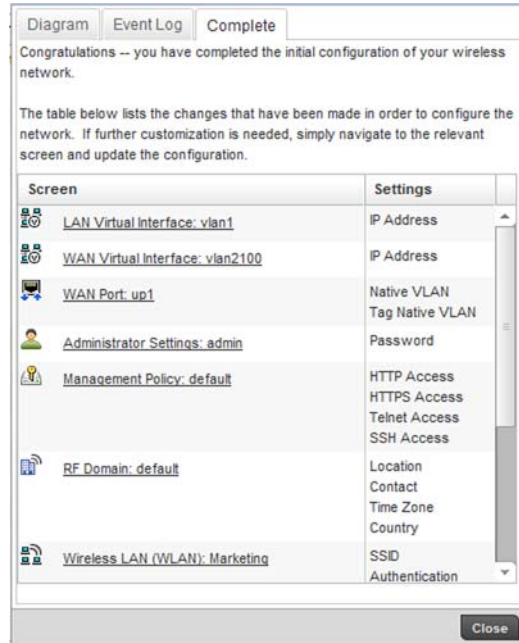


FIGURE 13 Initial Setup Wizard

The summary screen displays a table listing all changes made to the controller configuration by the wizard. It lists both the screens and the associated settings that have been modified.

Once you have reviewed the changes, select the **Close** button to exit the wizard and return the Web UI.

Creating a managed WLAN

This section describes the activities required to configure a managed WLAN. Instructions are provided using both the controller CLI and GUI to allow an administrator to configure the WLAN using the desired interface.

It's assumed you have a Brocade Mobility RFS4000 wireless controller with the latest build available from Brocade Mobility. It is also assumed you have one an Brocade Mobility 71XX Access Point model Access Point and one Brocade Mobility 650 Access Point model Access Point, both with the latest firmware from Brocade Mobility.

Upon completion, you'll have created a WLAN on a Brocade Mobility RFS4000 model wireless controller using a DHCP server to allocate IP addresses to associated wireless clients.

Assumptions

Creating a managed WLAN

Verify the following conditions have been satisfied before attempting the WLAN configuration activities described in this section.

It's assumed the wireless controller has the latest firmware version available from Brocade Mobility.

It's assumed the Brocade Mobility 71XX Access Point and Brocade Mobility 650 Access Point Access Points also have the latest firmware version available from Brocade Mobility.

It's assumed there are no previous configurations on the wireless controller or Access Point, and default factory configurations are running on the devices.

It's assumed you have administrative access to the RFS4000 wireless controller and Access Point GUI and CLI.

It's assumed the individual administrating the network is a professional network installer.

Design

[Creating a managed WLAN](#)

This section defines the network design being implemented.



FIGURE 14 Network Design

This is a fairly simple deployment scenario, with Access Points connected directly to the wireless controller. One wireless controller port is connected to an external network.

On the Brocade Mobility RFS4000 wireless controller, the GE1 interface is connected to an external network. Interfaces GE3 and GE4 are used by the access points.

On the external network, the controller is assigned an IP address of 192.168.10.188. The wireless controller acts as a DHCP server for the wireless clients connecting to it, and assigns IP addresses in the range of 172.16.11.11 to 172.16.11.200. The rest of IPs in the range are reserved for devices requiring static IP addresses.

To define the WLAN configuration using either controller GUI refer to:

- [Using the Controller GUI to Configure the WLAN](#)

Using the Controller GUI to Configure the WLAN

[Creating a managed WLAN](#)

The following instructions are for configuring a (non default) WLAN using the controller's *graphical user interface* (GUI).

Use a serial console cable when connecting to the wireless controller for the first time. Set the following configuration parameters when using a serial connection.

- Bits per second: *19200*
- Data Bit: *8*
- Parity: *None*
- Stop Bit: *1*
- Flow Control: *None*

When the wireless controller is started for the first time, its interfaces are not configured. Access to the wireless controller is only available through the serial console. To use the wireless controller's GUI, one of the other controller ports must be enabled and configured. The following section, demonstrates how to configure access to the GUI.

The tasks required to create a controller WLAN using the GUI include:

- [Configuring Access to the GUI Using the GE1 Port](#)
- [Logging into the Controller for the First Time](#)
- [Creating a RF Domain](#)
- [Creating a Wireless Controller Profile](#)
- [Creating a WLAN Configuration](#)
- [Creating an AP Profile](#)
- [Completing and testing the configurations](#)

Configuring Access to the GUI Using the GE1 Port

[Using the Controller GUI to Configure the WLAN](#)

Before you can access the wireless controller's GUI, the controller must have an IP address defined. The GE interface has to be configured with an IP address (using the CLI) before an administrator can access the GUI.

Logging into the Wireless Controller for the First Time

When you power on the wireless controller for the first time, you are prompted to replace the existing administrative password. The credentials for logging into the wireless controller for the first time include:

- User Name: *admin*
- Password: *admin123*

Ensure the new password created is strong enough to provide adequate security for the managed network.

Configuring the Controller's GE1 Interface

Navigate to the GE1 interface using the following commands:

```
Brocade Mobility RFS4000-571428>enable
```

```
Brocade Mobility RFS4000-571428#
```

```
Brocade Mobility RFS4000-571428#configure terminal
```

Enter configuration commands, one per line. End with CNTL/Z.


```
Brocade Mobility RFS4000-571428(config)#
```

```
Brocade Mobility RFS4000-571428(config)#self
```

```
Brocade Mobility RFS4000-571428(config-device-03-14-28-57-14-28)#
```

Create a VLAN and assign the IP address 172.16.0.1 to it.

```
Brocade Mobility RFS4000-571428(config-device-03-14-28-57-14-28)#interface vlan 1
```

```
Brocade Mobility RFS4000-571428(config-device-03-14-28-57-14-28-if-vlan1)#ip address 172.16.0.1
```

```
Brocade Mobility RFS4000-571428(config-device-03-14-28-57-14-28-if-vlan1)#commit write
```

```
Brocade Mobility RFS4000-571428(config-device-03-14-28-57-14-28)#
```

Configure the GE 1 port to use the VLAN 1.

```
Brocade Mobility RFS4000-571428(config-device-03-14-28-57-14-28)#interface ge 1
```

```
Brocade Mobility RFS4000-571428(config-device-03-14-28-57-14-28-if-ge1)#
```

```
Brocade Mobility RFS4000-571428(config-device-03-14-28-57-14-28-if-ge1)#switch port access vlan 1
```

The above command assigns the IP address 172.16.0.1, with the mask 255.255.255.0 to ME1. Exit the ME1 context.

```
Brocade Mobility RFS4000-571428(config-device-03-14-28-57-14-28-if-me1)#exit
```

```
Brocade Mobility RFS4000-571428(config-device-03-14-28-57-14-28)#
```

```
Brocade Mobility RFS4000-571428(config-device-03-14-28-57-14-28)#commit write
```

The system used to access the wireless controller must be configured as follows:

- *IP Address:* 172.16.0.10
- *Mask:* 255.255.255.0

Connect the device's Ethernet interface to the ME interface of the wireless controller.

Launch a browser and enter the following:

- <https://172.16.0.1/>

The controller's login screen displays.




FIGURE 15 Login Screen

Logging into the Controller for the First Time

Using the Controller GUI to Configure the WLAN

The following screen displays upon successfully changing the login password:

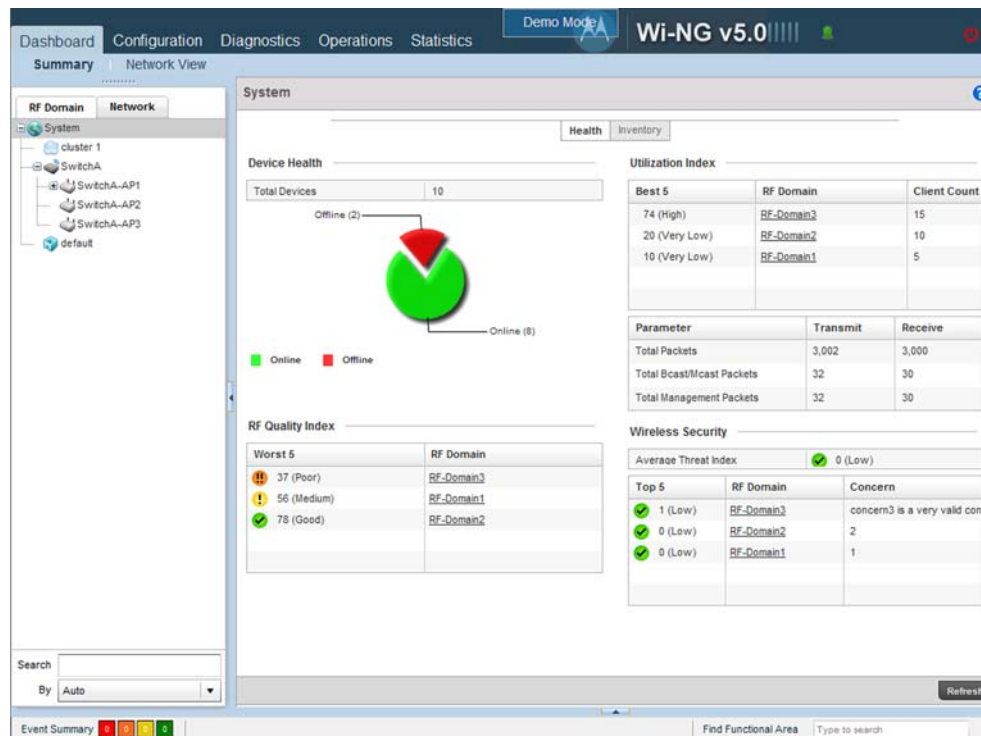


FIGURE 16 GUI Main screen

Creating a RF Domain

Using the Controller GUI to Configure the WLAN

A RF Domain is a collection of configuration settings specific to devices located at the same physical deployment, such as a building or a floor. Create a RF Domain and assign the country code where the devices are deployed. This is a mandatory step, and devices will not function as intended if this step is omitted.

To create a RF Domain:

Select **Configuration > RF Domain**.

RF Domain	Location	Contact	Time Zone	Country
default	test			United States-us

Type to Search in Tables

Row Count: 1

Add Edit Delete

FIGURE 17 RF Domain screen

Select the **Add** button at the bottom of the screen to create a new RF Domain configuration.

FIGURE 18 RF Domain screen - New RF Domain

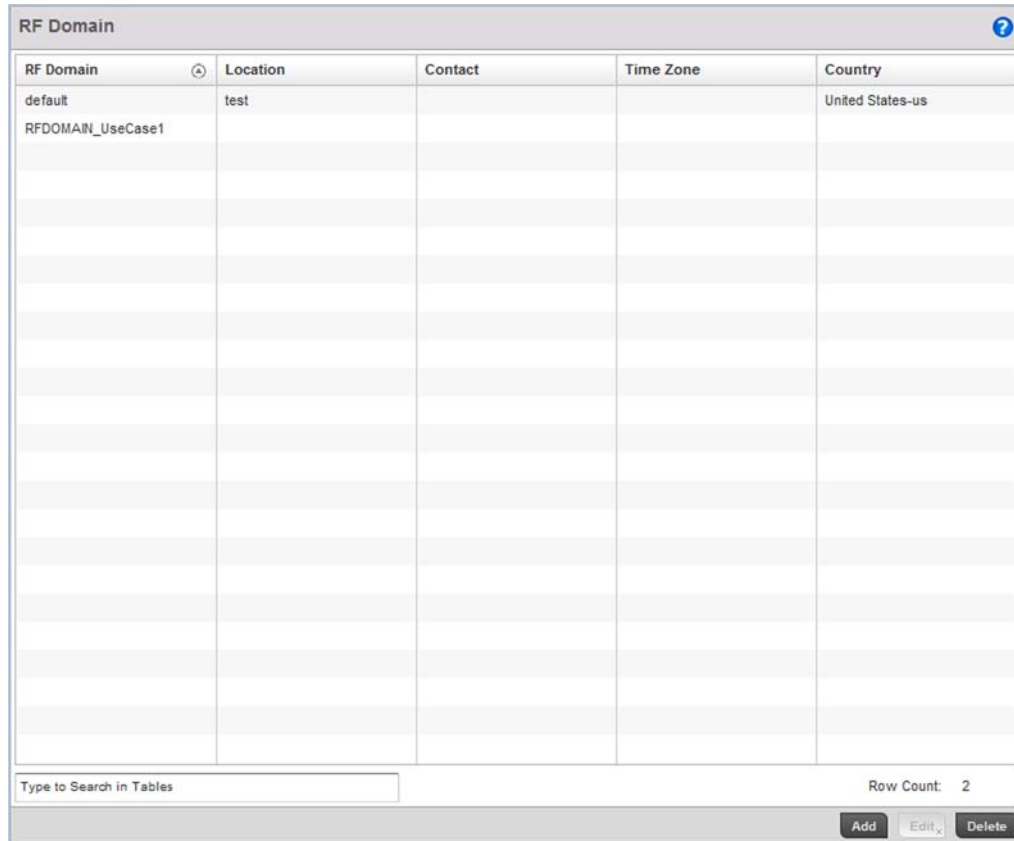
Provide the following information to define the RF Domain configuration:

RF Domain Assign the RF Domain a name representative of its intended function. The name cannot exceed 32 characters. The name cannot be changed as part of the edit process. For this scenario, use:
RFDOMAIN_UseCase1

Time Zone Assign the RF Domain a time zone representative of its deployment location. For this scenario, use:
(GMT - 08:00) America/Los_Angeles

Country Define the two-digit country code for the RF Domain. The country code must be set accurately to avoid the policy's illegal operation, as device radios transmit in specific channels unique to the country of operation. For this scenario, use:
United States - us

Select **OK** to save the updates. Select **Exit** button to close the screen, then click the **Commit** icon at the top right of the screen to apply the updates to the controller's running configuration.



The screenshot shows a web-based configuration interface for RF Domains. At the top, there is a header bar labeled "RF Domain" with a help icon on the right. Below the header is a table with five columns: "RF Domain", "Location", "Contact", "Time Zone", and "Country". The table contains two rows of data. The first row has "default" in the "RF Domain" column, "test" in the "Location" column, and "United States-us" in the "Country" column. The second row has "RFDOMAIN_UseCase1" in the "RF Domain" column. Below the table, there is a search input field labeled "Type to Search in Tables" and a "Row Count: 2" indicator. At the bottom right, there are three buttons: "Add", "Edit", and "Delete".

RF Domain	Location	Contact	Time Zone	Country
default	test			United States-us
RFDOMAIN_UseCase1				

FIGURE 19 New RF Domain Configuration

Configure the Wireless Controller to use the RF Domain

To configure the wireless controller's physical deployment location, use the RF Domain configuration you just created.

Select **Configuration > Devices > Device Configuration**.

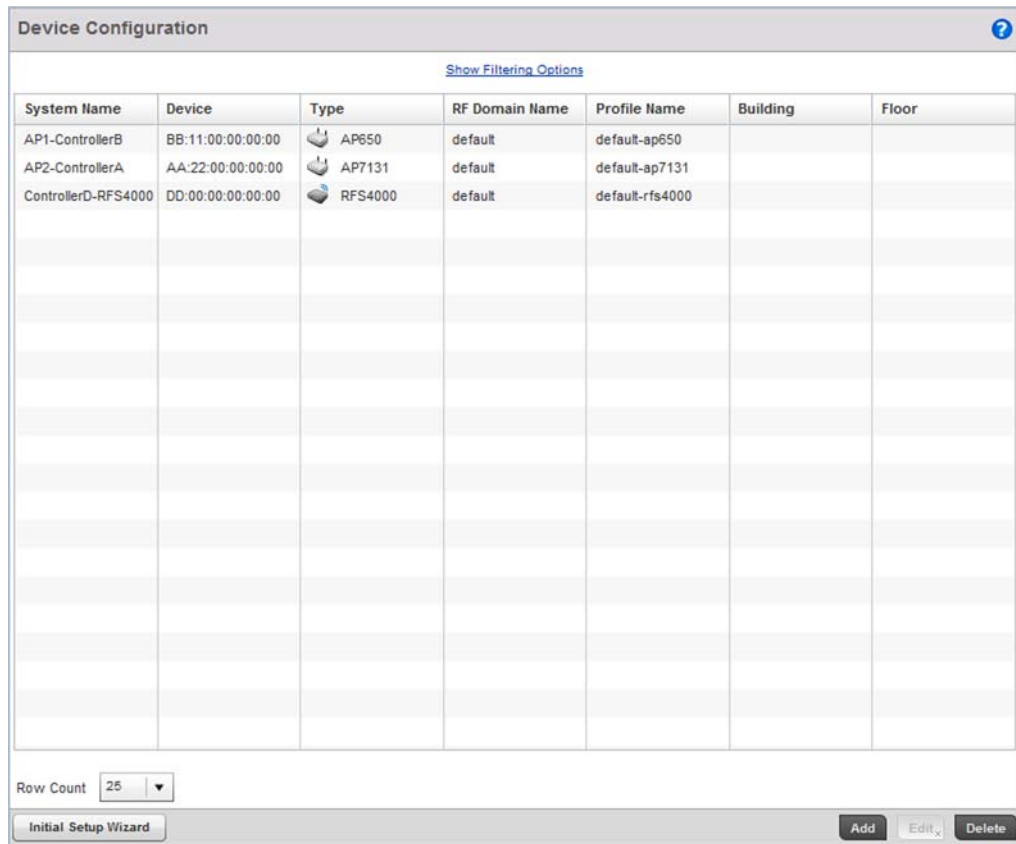


FIGURE 20 Device Configuration screen

Select the Brocade Mobility RFS4000 wireless controller. Select the **Edit** button.

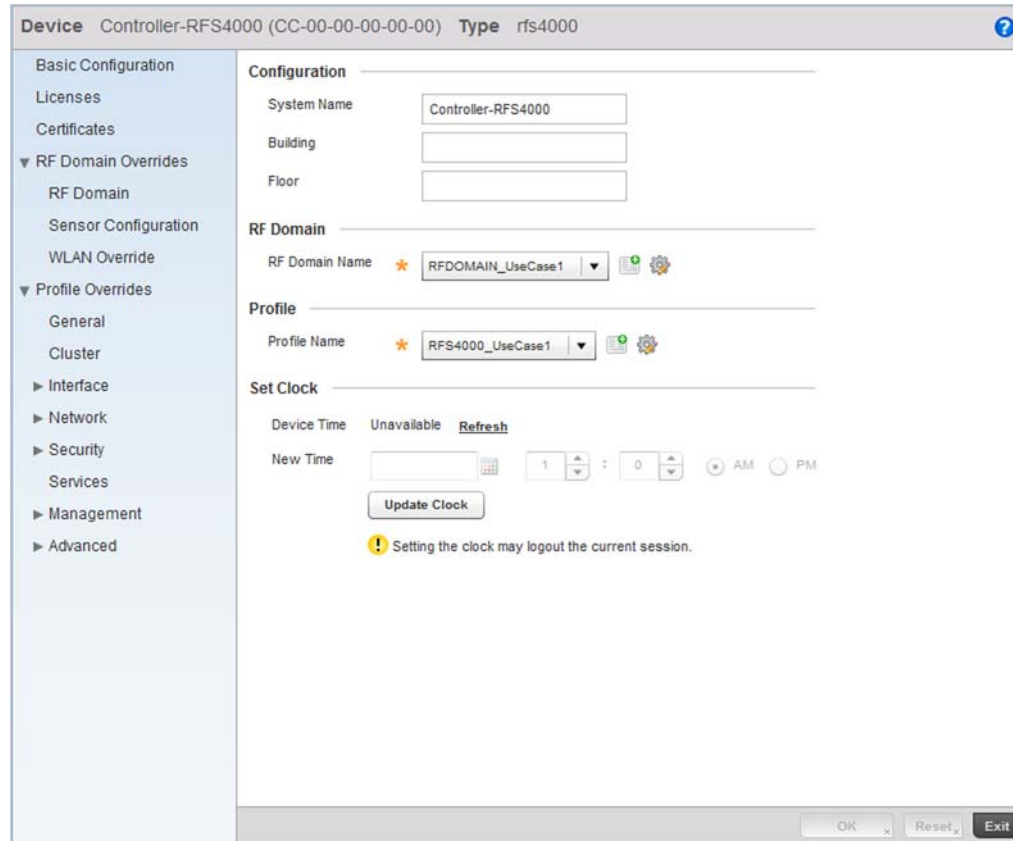


FIGURE 21 Brocade Mobility RFS4000 Device screen

Set a name for the RF Domain to which this Brocade Mobility RFS4000 controller belongs. For this use case scenario, use:

RFDOMAIN_UseCase1

Leave the rest of the fields undefined.

Select **OK** button to save the changes. Select **Exit** button to close the screen, then click the **Commit** icon at the top right of the screen to apply the updates to the controller's running configuration.

Repeat the steps 1 through 4 to configure the Brocade Mobility 650 Access Point.

NOTE

The wireless controller and Access Points must use the same country code for the Access Points to be successfully adopted by the controller.

The following image displays after the Brocade Mobility RFS4000, Brocade Mobility 650 Access Point and Brocade Mobility 71XX Access Point have been configured to use the RFDOMAIN_UseCase1 RF Domain.

System Name	Device	Type	RF Domain Name	Profile Name	Building	Floor
AP1-ControllerA	AA:11:00:00:00:00	AP650	RFDOMAIN_UseCase1	AP650_UseCase		
AP2-ControllerA	AA:22:00:00:00:00	AP7131	RFDOMAIN_UseCase1	AP7131_UseCase1		
Controller-RFS4000	CC:00:00:00:00:00	RFS4000	RFDOMAIN_UseCase1	RFS4000_UseCase1		

Row Count 25

Initial Setup Wizard Add Edit Delete

FIGURE 22 Device Configuration screen after configuring Brocade Mobility RFS4000 and Brocade Mobility 650 Access Point

Creating a Wireless Controller Profile

Using the Controller GUI to Configure the WLAN

The first step in creating a WLAN is to create a profile defining the configuration parameters that must be applied to the wireless controller.

To create a profile:

Select **Configuration > Profiles**.

Profile ?								
Profile ⌵	Type	Adoption Policy	Firewall Policy	Wireless Client Role Policy	Advanced WIPS Policy	DHCP Server Policy	Management Policy	RADIUS Server Policy
default-ap650	AP650							
default-ap7131	AP7131						management1	
default-rfs4000	RFS4000							

Row Count: 3

Add Edit Delete

FIGURE 23 Profiles screen

Select the **Add** button at the bottom right of the screen.

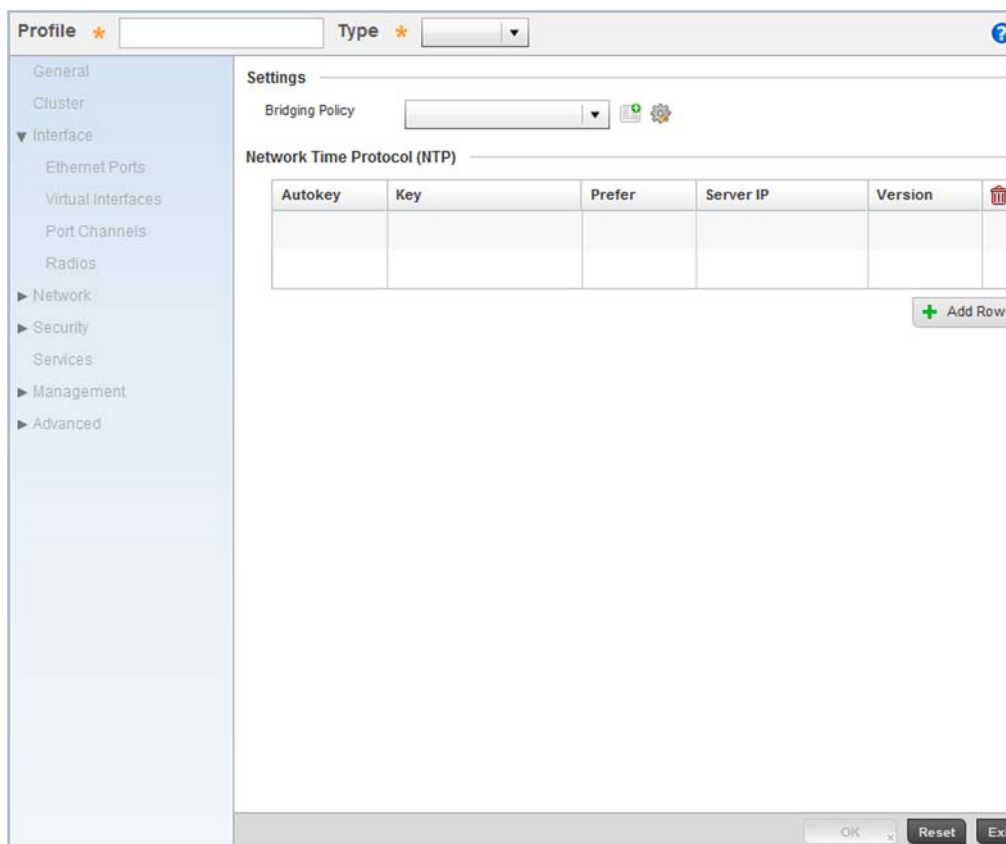


FIGURE 24 New Brocade Mobility RFS4000 Profile

Define the name of the profile and the device type it supports.

Profile	Define the name of the Brocade Mobility RFS4000 profile being created. For this scenario, use: Brocade Mobility RFS4000_UseCase1
Type	Specify the device type. For this scenario, use: Brocade Mobility RFS4000

Select **OK** to save the changes. Select the **Exit** button to close this screen, then click the **Commit** icon at the top right of the screen to apply the updates to the controller's running configuration.

This creates a profile with the name *Brocade Mobility RFS4000_UseCase1*. Any configuration made under this profile is available for use when it's applied to a device.

Configure a VLAN

Create the VLAN used with this WLAN.

Using the Brocade Mobility RFS4000_UseCase1 profile, expand the **Configuration > Profiles > Interface** menu item to display its submenu options.

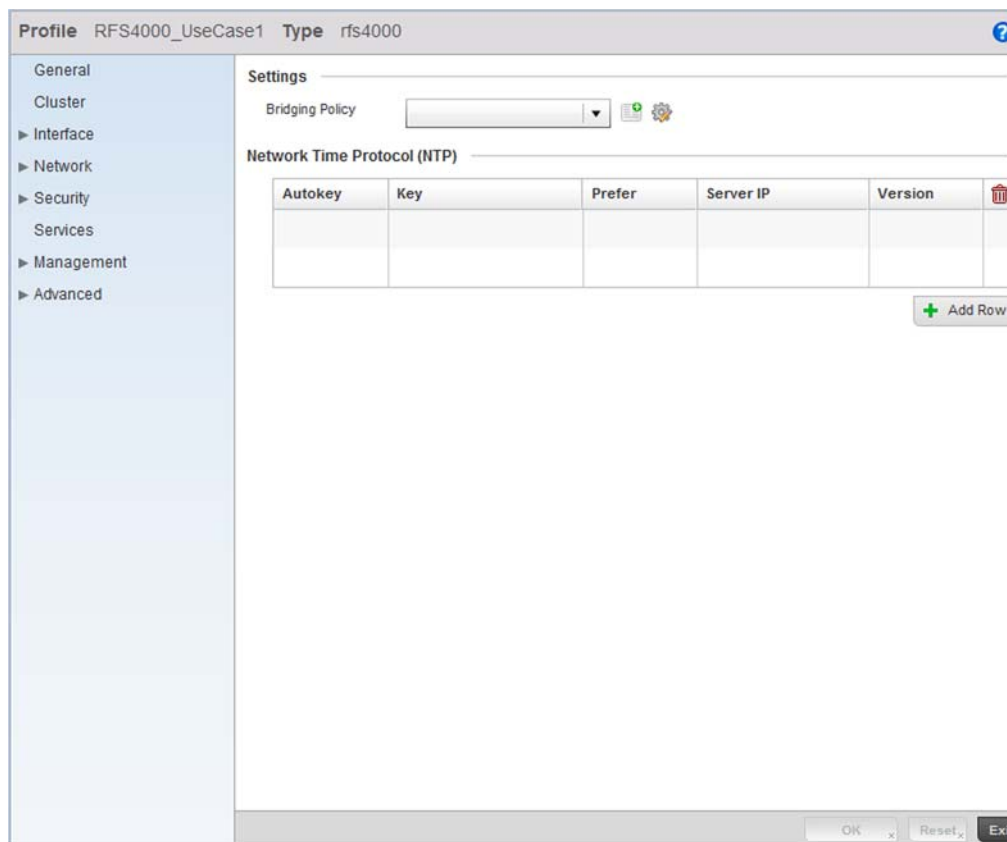


FIGURE 25 Brocade Mobility RFS4000 Profile screen

Select **Virtual Interfaces**.

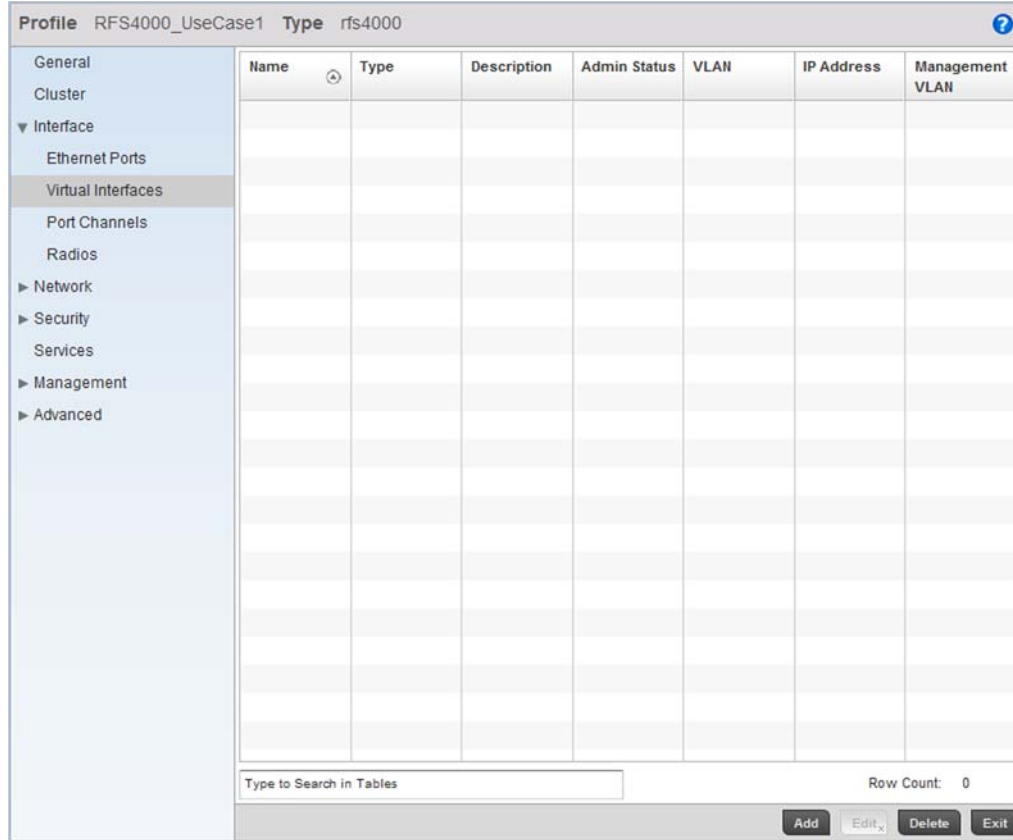


FIGURE 26 Virtual Interfaces screen

Click the **Add** button located at the bottom left of the screen.

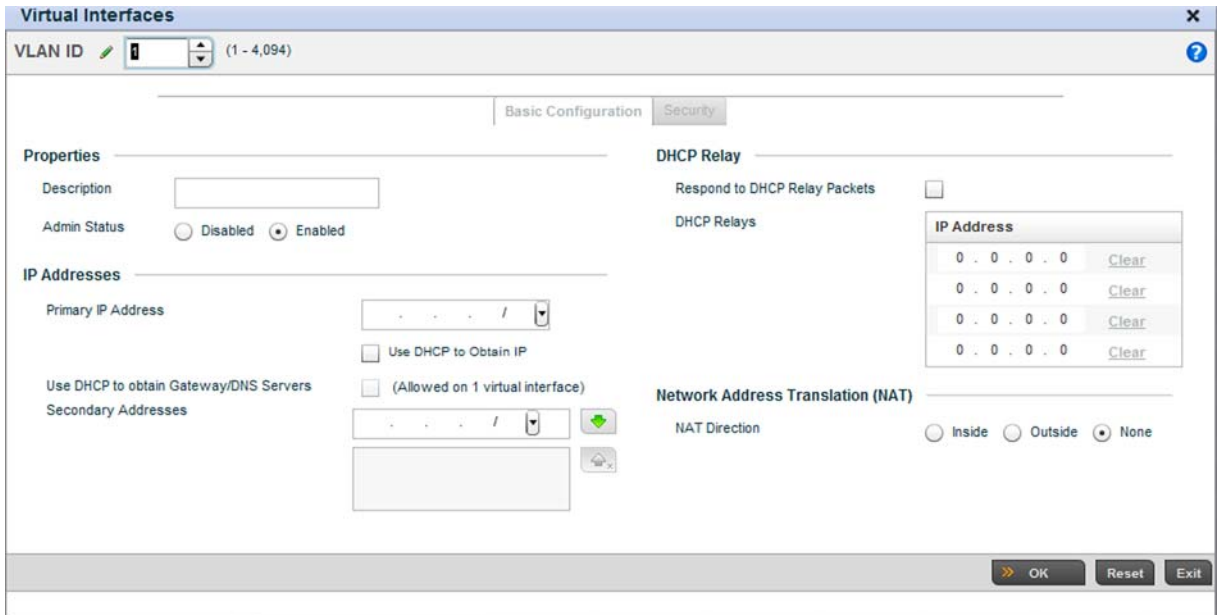


FIGURE 27 Virtual Interface screen

Set a **VLAN ID** (within the top of the screen) and a **Primary IP Address**. For this use case scenario, use a VLAN ID of 2 and a Primary IP Address of 172.16.11.1/24. This assigns an IP address of 172.16.11.1 with a mask of 255.255.255.0 to VLAN2.

Select **OK** to save the changes. Select **Exit** to close the screen, then click the **Commit** icon at the top right of the screen to apply the updates to the controller's running configuration.

The next step is to assign this newly created VLAN to a physical interface. In this case, VLAN 2 is mapped to GE3 and GE4 to support the Brocade Mobility 650 Access Point and an Brocade Mobility 71XX Access Point Access Points. The Brocade Mobility 650 Access Point is connected to the gigabit interface GE3, and the Brocade Mobility 71XX Access Point to the interface GE4.

Configure the Physical Interfaces

To configure the GE3 port on behalf of the Brocade Mobility 650 Access Point:

Using the RFS7000_UseCase1 profile, expand the **Configuration > Profiles > Interface** menu item to display its submenu options.

Select **Ethernet Ports**.

Name	Type	Description	Admin Status	Mode	Native VLAN	Tag Native VLAN	Allowed VLANs
ge1	Ethernet		✓ Enabled	Access	1	✗	
ge2	Ethernet		✓ Enabled	Access	1	✗	
ge3	Ethernet		✓ Enabled	Access	1	✗	
ge4	Ethernet		✓ Enabled	Access	1	✗	
ge5	Ethernet		✓ Enabled	Access	1	✗	
up1	Ethernet		✓ Enabled	Access	1	✗	

FIGURE 28 Ethernet Port Configuration screen

By default, all ports are enabled and VLAN 1 is assigned as the Native VLAN. To change the Native VLAN value to VLAN 2, select the GE 3 controller interface and select the **Edit** button.

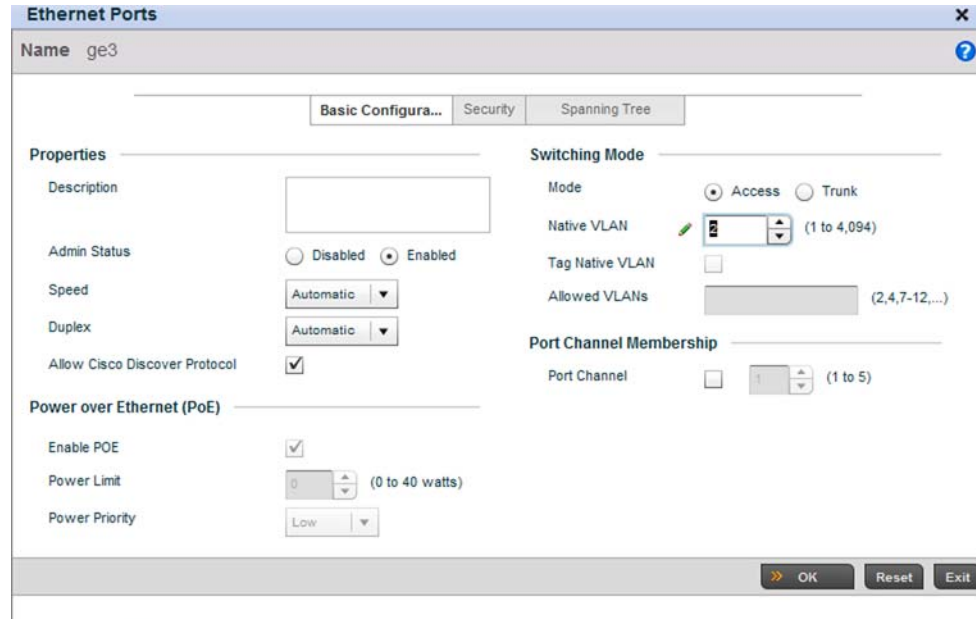


FIGURE 29 Interface GE 3 Configuration

Use the spinner control to set a **Native VLAN** value of 2 for this use case scenario.

Do not change any other values.

Click the **OK** button to save the changes. Select **Exit** to close the screen, then click the **Commit** icon at the top right of the screen to apply the updates to the controller's running configuration.

Repeat the steps 3 through 5 to map VLAN 2 to the GE 4 interface. Once completed, the screen should appear like [Figure 30](#).

Profile default-rfs4000 Type rfs4000

Name	Type	Description	Admin Status	Mode	Native VLAN	Tag Native VLAN	Allowed VLANs
ge1	Ethernet		✓ Enabled	Access	1	✗	
ge2	Ethernet		✓ Enabled	Access	1	✗	
ge3	Ethernet		✓ Enabled	Access	2	✗	
ge4	Ethernet		✓ Enabled	Access	2	✗	
ge5	Ethernet		✓ Enabled	Access	1	✗	
up1	Ethernet		✓ Enabled	Access	1	✗	

Type to Search in Tables Row Count: 6

FIGURE 30 Physical Interfaces GE3 and GE4




Configuring the Wireless Controller to use the Appropriate Profile

Before the wireless controller can be configured further, the profile must be applied to the wireless controller. To do so:

Navigate to the Brocade Mobility RFS4000 by selecting **Configuration > Devices > Device Configuration**.

Device Configuration ?

[Show Filtering Options](#)

System Name	Device	Type	RF Domain Name	Profile Name	Building	Floor
AP1-ControllerA	AA:11:00:00:00:00	 AP650	RFDOMAIN_UseCase1	AP650_UseCase		
AP2-ControllerA	AA:22:00:00:00:00	 AP7131	RFDOMAIN_UseCase1	AP7131_UseCase1		
Controller-RFS4000	CC:00:00:00:00:00	 RFS4000	RFDOMAIN_UseCase1	RFS4000_UseCase1		

Row Count ▼

Initial Setup Wizard Add Edit Delete

FIGURE 31 Device Configuration screen

Select the Brocade Mobility RFS4000 from the list and select the **Edit** button.

FIGURE 32 Brocade Mobility RFS4000 Configuration

Provide a profile name for the Brocade Mobility RFS4000 device profile. For this use case scenario, use:

Brocade Mobility RFS4000_UseCase1

Select **OK** to save the changes. Select **Exit** to close this screen, then click the **Exit** button to close this screen. Select the **Commit** icon to save it to the controller's running configuration.

Creating a WLAN Configuration

Using the Controller GUI to Configure the WLAN

Complete the following steps to create a WLAN:

Select **Configuration > Wireless > Wireless LANs** to navigate to the WLAN screen.

Wireless LANs ?

WLAN ⊙	SSID	Description	WLAN Status	VLAN Pool	Authentication Type	Encryption Type	QoS Policy	Association ACL

Type to Search in Tables Row Count: 0

Add **Edit** **Delete**

FIGURE 33 Wireless LANs screen

Select the **Add** button to create a new WLAN.

FIGURE 34 WLAN Configuration screen

Provide the following information to define the controller WLAN configuration.

- WLAN** Define the name of the WLAN. For this scenario, use 1.
- SSID** Provide the *Service Set Identifier* (SSID) for the WLAN. This is the ID used when Access Points and wireless clients need to associate with the WLAN. For this scenario, use: WLAN_USECASE_01
- VLAN** Define the VLAN to associate with WLAN_USECASE_01. For this scenario, use: VLAN 2
- Single VLAN** Ensure this option is selected to restrict WLAN_USECASE_01's VLAN usage to VLAN 2.

Select **OK** to save the changes. Select **Exit** to close this screen, then select the **Commit** icon to save it to the controller running configuration. Once completed, the screen should appear as [Figure 35](#).

Profile	Type	Adoption Policy	Firewall Policy	Wireless Client Role Policy	Advanced WIPS Policy	DHCP Server Policy	Management Policy	RADIUS Server Policy
default-ap650	AP650							
default-ap7131	AP7131						management1	
default-rfs4000	RFS4000							
RFS4000_UseC	RFS4000							

Type to Search in Tables Row Count: 4

Add **Edit** **Delete**

FIGURE 36 Profiles screen

Select the **Add** button located at the bottom right of the screen.

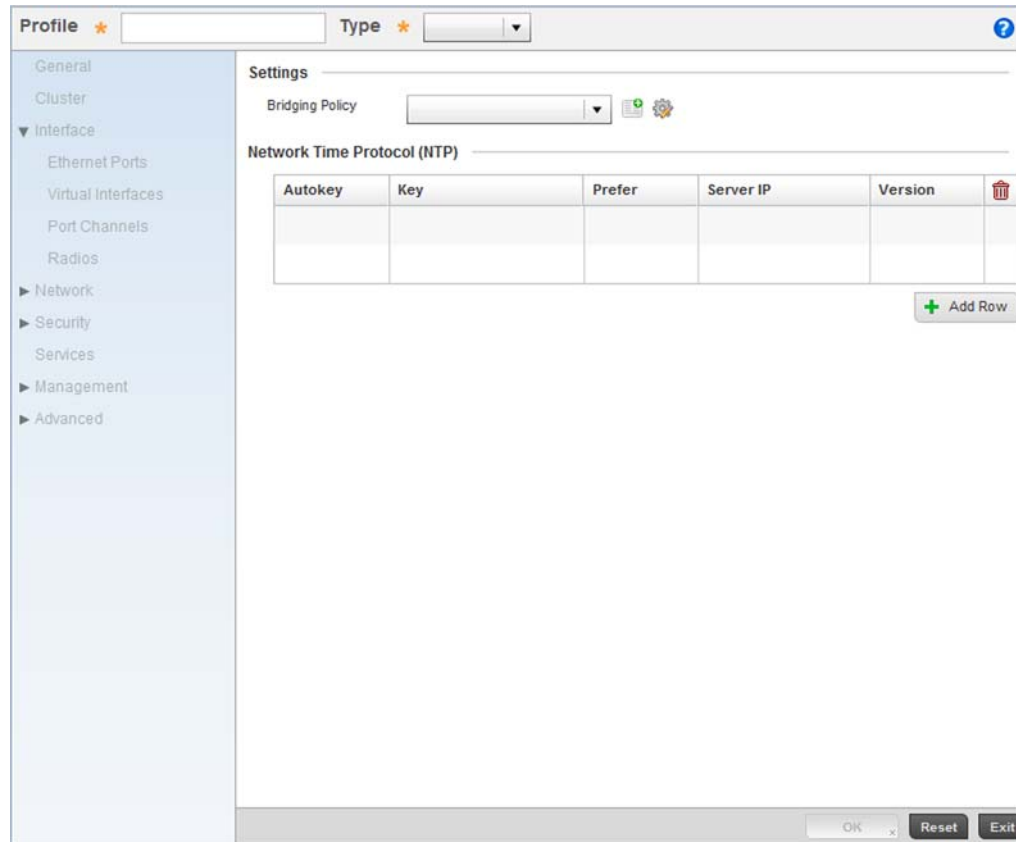


FIGURE 37 New Brocade Mobility 650 Access Point Profile

Provide a **Profile Name** for the new Brocade Mobility 650 Access Point profile. For this scenario, use: Brocade Mobility 650 Access Point_UseCase1.

Define the device **Type** for this profile. For this use case scenario, use: Brocade Mobility 650 Access Point.

Select **OK** to save the changes.

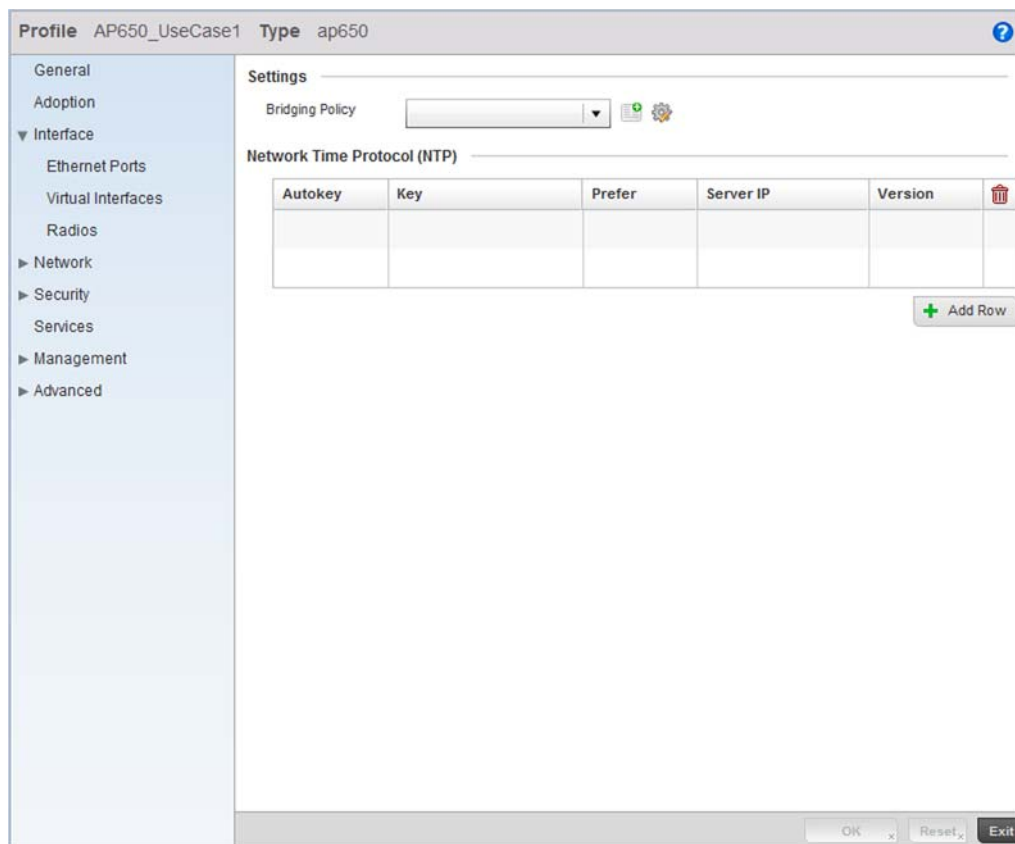


FIGURE 38 Brocade Mobility 650 Access Point Profile

Select the **Interface** menu item to expand it and display its submenu items.

Select **Virtual Interfaces**.

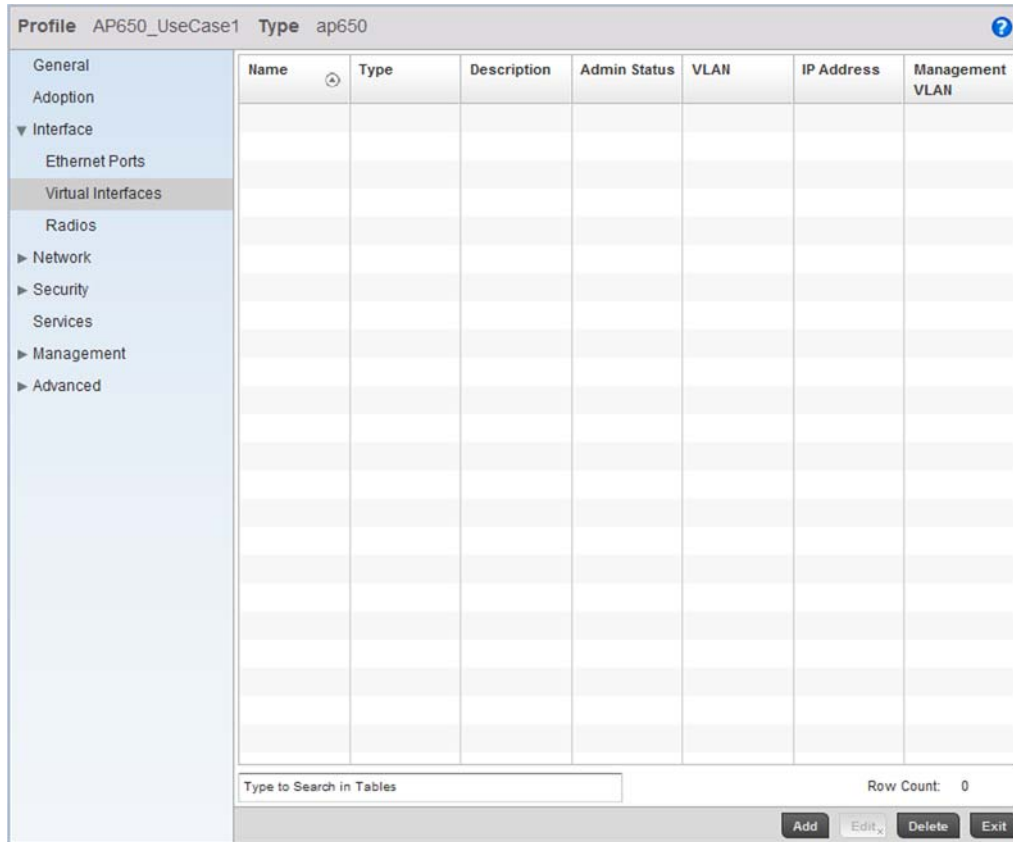


FIGURE 39 Configuring Virtual Interfaces for an Brocade Mobility 650 Access Point profile.

Select the **Add** button at the bottom of the screen. A screen displays for adding a new virtual interface configuration for the profile.

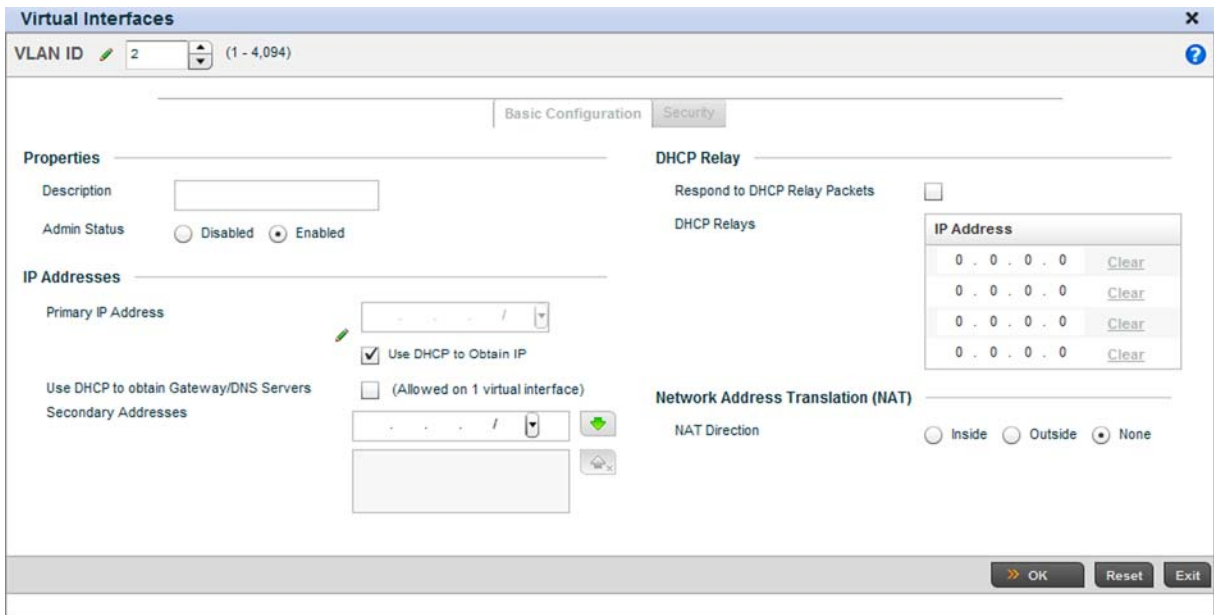


FIGURE 40 Virtual Interface screen

Define a **VLAN ID** for the Brocade Mobility 650 Access Point profile. For this scenario, set it as VLAN 2.

For the **Primary IP Address** assignment, select the **Use DHCP to Obtain IP** option.

Select the **OK** button to save the changes. Click the **Exit** button to close the screen, then click the **Commit** icon at the top right of the screen to apply the updates to the controller's running configuration.

Configure the Physical Interface

The next step is to map this newly created VLAN to a physical interface. In this case, VLAN 2 is mapped to the GE1 port of the Brocade Mobility 650 Access Point.

To configure the GE1 port:

Select and display the Brocade Mobility 650 Access Point_UseCase1 profile by navigating to **Configuration > Profiles** and selecting the profile from amongst those displayed.

Autokey	Key	Prefer	Server IP	Version	

FIGURE 41 Brocade Mobility 650 Access Point Profile screen

From within the Profiles screen, select and expand the **Interface** menu to display its submenu items.

Select **Ethernet Ports**.

Ethernet Ports ✕

Name `ge1` ?

Basic Configuration Security Spanning Tree

Properties

Description

Admin Status Disabled Enabled

Speed ▼

Duplex ▼

Allow Cisco Discover Protocol

Power over Ethernet (PoE)

Enable POE

Power Limit (0 to 40 watts)

Power Priority ▼

Switching Mode

Mode Access Trunk

Native VLAN (1 to 4,094)

Tag Native VLAN

Allowed VLANs (2,4,7-12,...)

Port Channel Membership

Port Channel (1 to 5)

FIGURE 42 Interface GE1 Configuration screen

Use the spinner control to define the **Native VLAN**. For this scenario, select 2. No other value on the screen requires configuration.

Select **OK** to save the changes. Select **Exit** to close the screen, then click the **Commit** icon at the top right of the screen to apply the updates to the controller's running configuration. Once completed, the Profiles screen for the Brocade Mobility 650 Access Point should appear as [Figure 43](#).

Name	Type	Description	Admin Status	Mode	Native VLAN	Tag Native VLAN	Allowed VLANs
ge1	Ethernet		✓ Enabled	Access	2	✗	

FIGURE 43 Brocade Mobility 650 Access Point Profiles screen after configuring the GE1 Physical Interface

Configure the Brocade Mobility 650 Access Point Radios

Each configured WLAN must be assigned an Access Point radio before wireless clients can connect to it. To configure the Brocade Mobility 650 Access Point's radios:

Select and display the Brocade Mobility 650 Access Point_UseCase1 profile by navigating to **Configuration > Profiles** and selecting the profile from amongst those displayed.

From within the Profiles screen, select and expand the **Interface** menu to display its submenu items.

Select **Radios**.

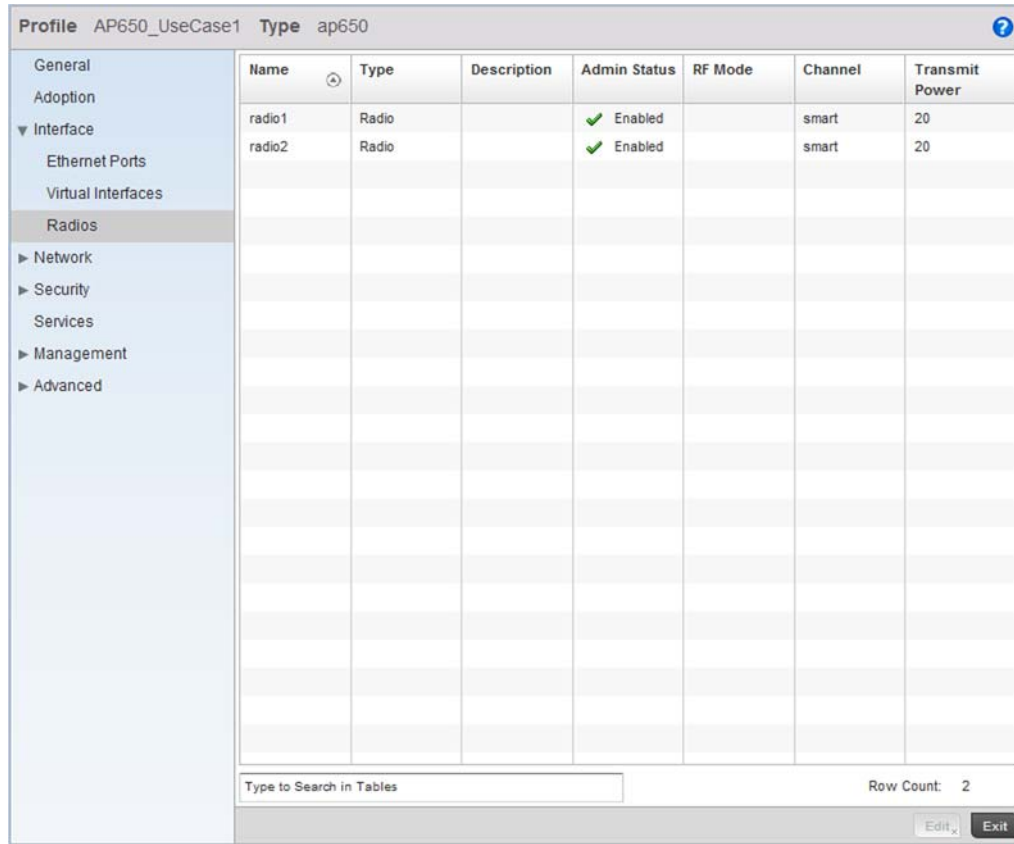


FIGURE 44 Radios Screen

Select **Radio1**, then select the **Edit** button.

The screenshot shows the 'Radio 1' configuration window with the following sections and settings:

- Radio Settings** (Active Tab):
 - Radio QoS Policy: Disabled, Enabled
 - Association ACL:
 - RF Mode:
 - Lock Radio Band:
 - Channel: smart
 - Transmit Power: smart, 20
 - Antenna Gain: 0.00 (0.00 - 15.00 dBi)
 - Antenna Mode: Default
 - Dynamic Chain Selection:
 - Data Rates: Select
 - Radio Placement: Indoor
- WLAN Properties**:
 - Beacon Interval: 100 (milliseconds)
 - DTIM Interval BSSID: Select
 - RTS Threshold: 2347 (1 to 2,347 bytes)
 - Short Preamble:
 - Guard Interval: Any
 - Probe Response Rate: follow-probe-request
 - Probe Response Retry:
- Channel Scanning**:
 - Enable Off Channel Scan:

Buttons at the bottom: OK, Reset, Exit.

FIGURE 45 Radio 1 Configuration screen

Select the **WLAN Mapping** tab.

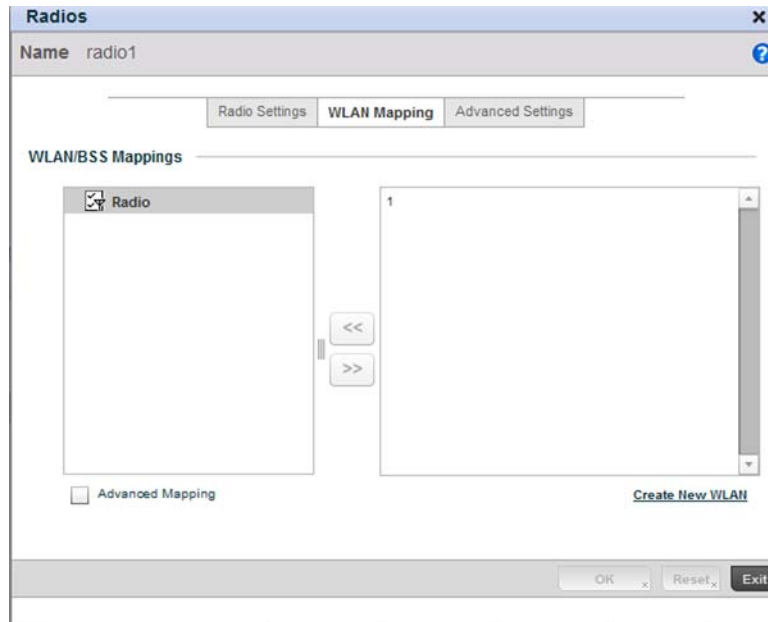


FIGURE 46 Radio 1 Configuration - WLAN Mapping screen.

From the list on the right of the screen, select the WLAN to assign to this radio. Select the << button to assign the selected WLAN to Radio 1. The screen updates as displayed in Figure 47.

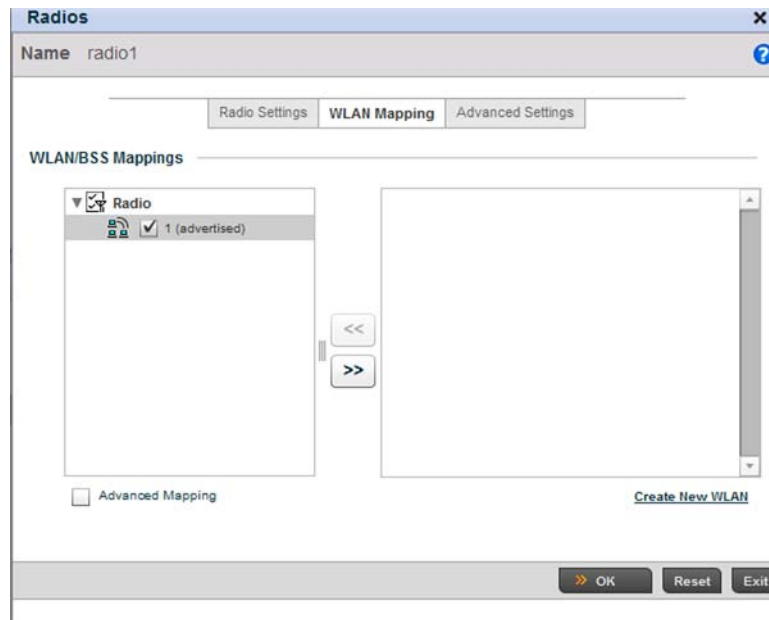


FIGURE 47 WLAN assigned to Radio1

Click the **OK** button to save the changes. Select **Exit** to exit the screen, then click the **Commit** button to write this change to the configuration.

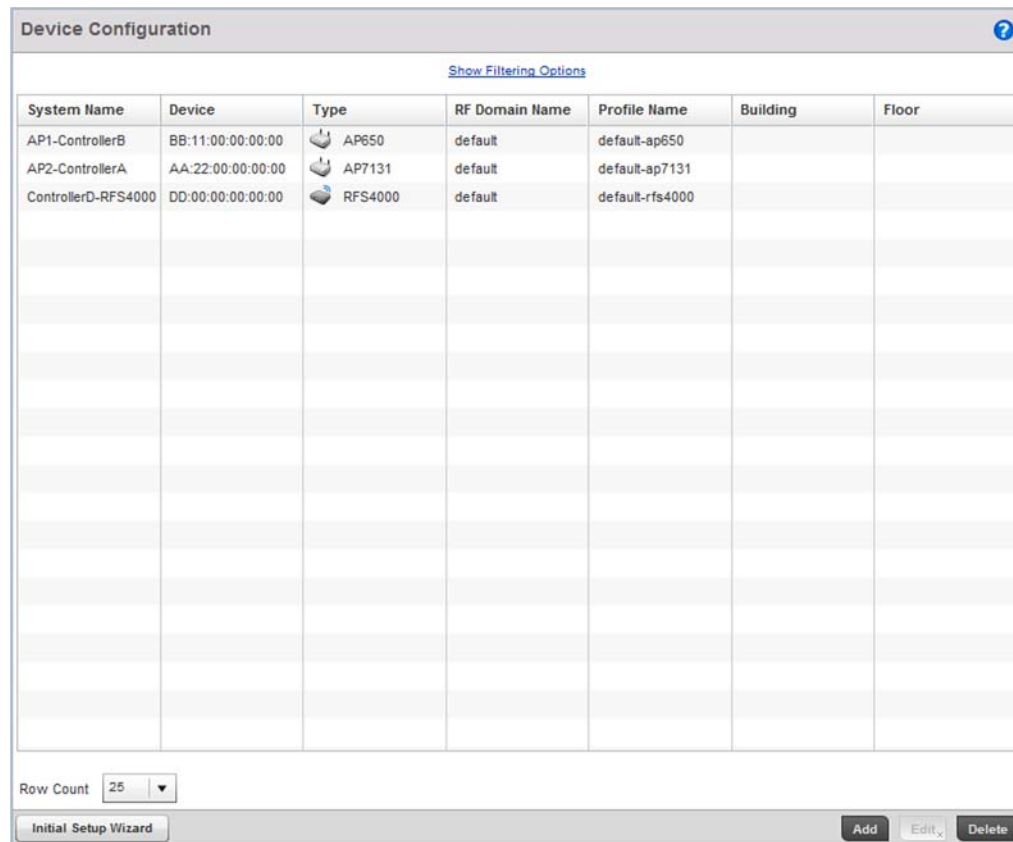
Configure the Brocade Mobility 650 Access Point's Radio 2 by repeating steps 3 through 7.

Configure the Brocade Mobility 650 Access Point to use the Profile

Before the Brocade Mobility 650 Access Point can be used as a managed device, the profile must be applied to the AP.

To apply the profile to the Access Point:

Navigate to the Brocade Mobility 650 Access Point by selecting **Configuration > Devices > Device Configuration**.



The screenshot shows the 'Device Configuration' screen with a table of devices. The table has columns for System Name, Device, Type, RF Domain Name, Profile Name, Building, and Floor. There are three rows of data, each with a small icon next to the Type column.

System Name	Device	Type	RF Domain Name	Profile Name	Building	Floor
AP1-ControllerB	BB:11:00:00:00:00	AP650	default	default-ap650		
AP2-ControllerA	AA:22:00:00:00:00	AP7131	default	default-ap7131		
ControllerD-RFS4000	DD:00:00:00:00:00	RFS4000	default	default-rfs4000		

At the bottom of the screen, there is a 'Row Count' dropdown set to 25, an 'Initial Setup Wizard' button, and 'Add', 'Edit', and 'Delete' buttons.

FIGURE 48 Device Configuration screen

Select Brocade Mobility 650 Access Point from amongst the devices displayed and select the **Edit** button.

The screenshot displays the configuration page for a Brocade Mobility 650 Access Point. The top header shows the device name 'AP1-ControllerB (BB-11-00-00-00-00)' and type 'ap650'. The left sidebar lists various configuration categories, with 'Profile Overrides' expanded to show 'General', 'Adoption', 'Interface', 'Network', 'Security', 'Services', 'Management', and 'Advanced'. The main configuration area is divided into sections: 'Configuration' with fields for System Name, Building, and Floor; 'RF Domain' with a dropdown for RF Domain Name; 'Profile' with a dropdown for Profile Name; and 'Set Clock' with a 'Device Time' status, a 'New Time' selector, and an 'Update Clock' button. A yellow warning icon and message are present below the clock settings. At the bottom right, there are 'OK', 'Reset', and 'Exit' buttons.

FIGURE 49 Brocade Mobility 650 Access Point Configuration

Define a **Profile Name** to use with this Brocade Mobility 650 Access Point. This applies the properties defined in the profile to the selected Brocade Mobility 650 Access Point. For this use case scenario, use Brocade Mobility 650 Access Point_UseCase1.

Select **OK** to save the changes. Select **Exit** to close the screen, then click the **Commit** icon at the top right of the screen to apply the updates to the controller's running configuration.

Creating an Brocade Mobility 71XX Access Point Profile

Creating an AP Profile

An Brocade Mobility 71XX Access Point is a standalone access point that provides small and medium businesses a cost effective device that consolidates a wired and wireless network infrastructure in a single device. It integrates a router, gateway, firewall and other services to simplify and reduce the overall cost of ownership by eliminating the need to maintain multiple devices.

To create an Brocade Mobility 71XX Access Point profile:

Navigate to the Profiles screen by selecting **Configuration > Profiles > Manage Profiles**.

Profile ?								
Profile ⊕	Type	Adoption Policy	Firewall Policy	Wireless Client Role Policy	Advanced WIPS Policy	DHCP Server Policy	Management Policy	RADIUS Server Policy
AP650_UseCase1	AP650							
default-ap650	AP650							
default-ap7131	AP7131						management1	
default-rfs4000	RFS4000							
RFS4000_UseCase1	RFS4000							

Type to Search in Tables Row Count: 5

[Add](#) [Edit](#) [Delete](#)

FIGURE 50 Profiles screen

To create a new profile, select the **Add** button at the bottom right of the screen.

The screenshot displays the configuration interface for a new Brocade Mobility 71XX Access Point profile. The interface includes a left-hand navigation pane with categories like General, Cluster, Interface, Network, Security, Services, Management, and Advanced. The main area is titled 'Settings' and contains a 'Bridging Policy' dropdown menu. Below this is a section for 'Network Time Protocol (NTP)' which features a table with the following columns: Autokey, Key, Prefer, Server IP, and Version. A '+ Add Row' button is positioned to the right of the table. At the bottom right of the window, there are three buttons: 'OK', 'Reset', and 'Exit'.

FIGURE 51 New Brocade Mobility 71XX Access Point Profile Creation

Define a name for the new AP7131 profile. For the purposes of this use case scenario, use Brocade Mobility 71XX Access Point_UseCase1.

Ensure the device type is set as AP7131 from the **Type** drop-down menu.

Select **OK** to save the changes.

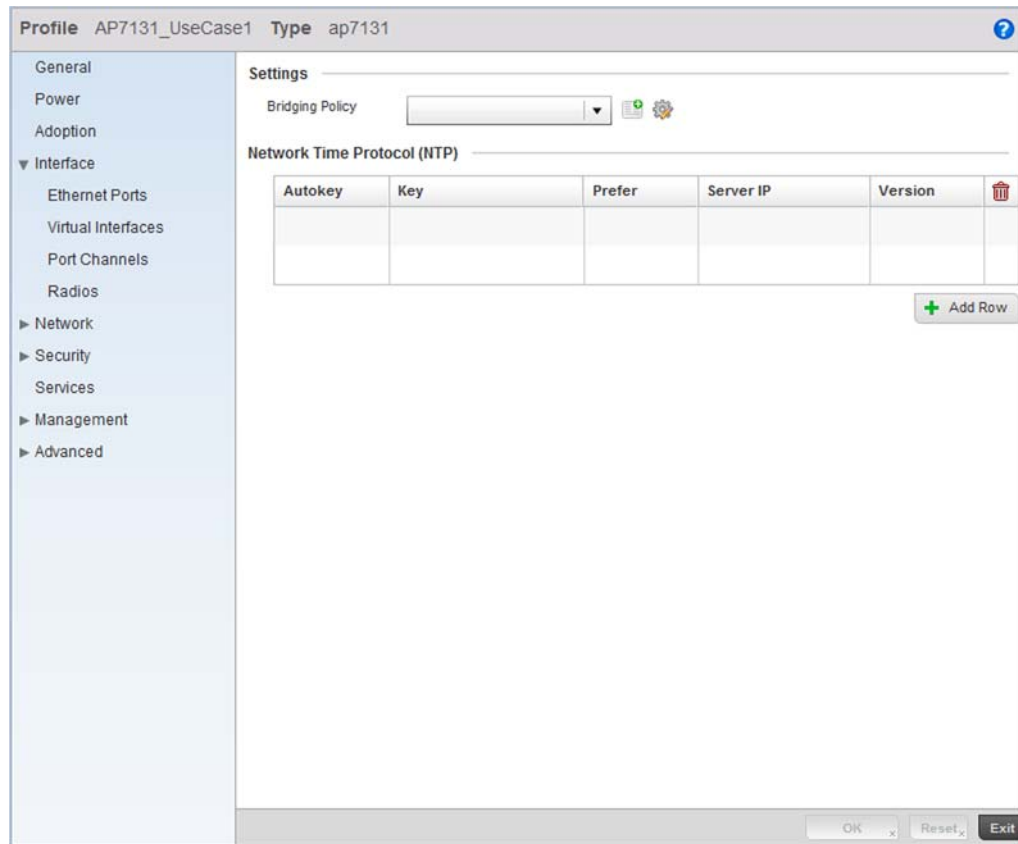


FIGURE 52 Brocade Mobility 71XX Access Point Profile

Select the **Interface** menu item to expand it and display its sub menu options.

Select **Virtual Interfaces**.

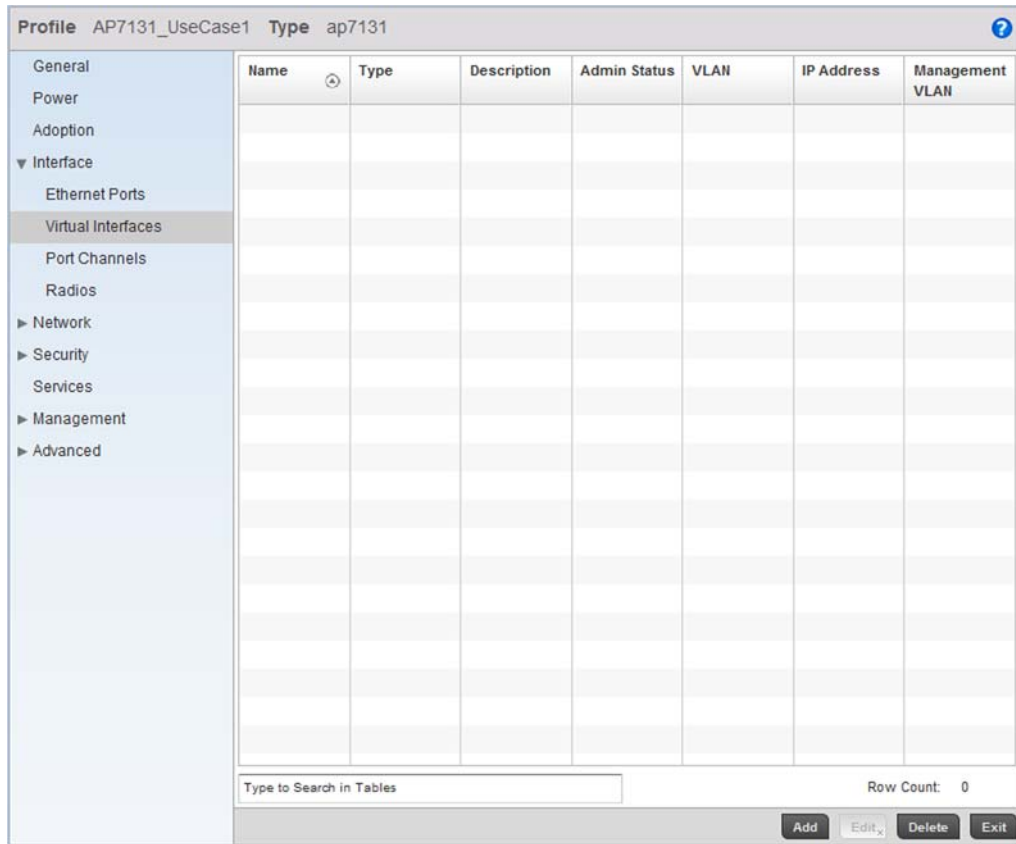


FIGURE 53 Configuring a Virtual Interface for the Brocade Mobility 71XX Access Point Profile.

Select the **Add** button at the bottom left of the screen.

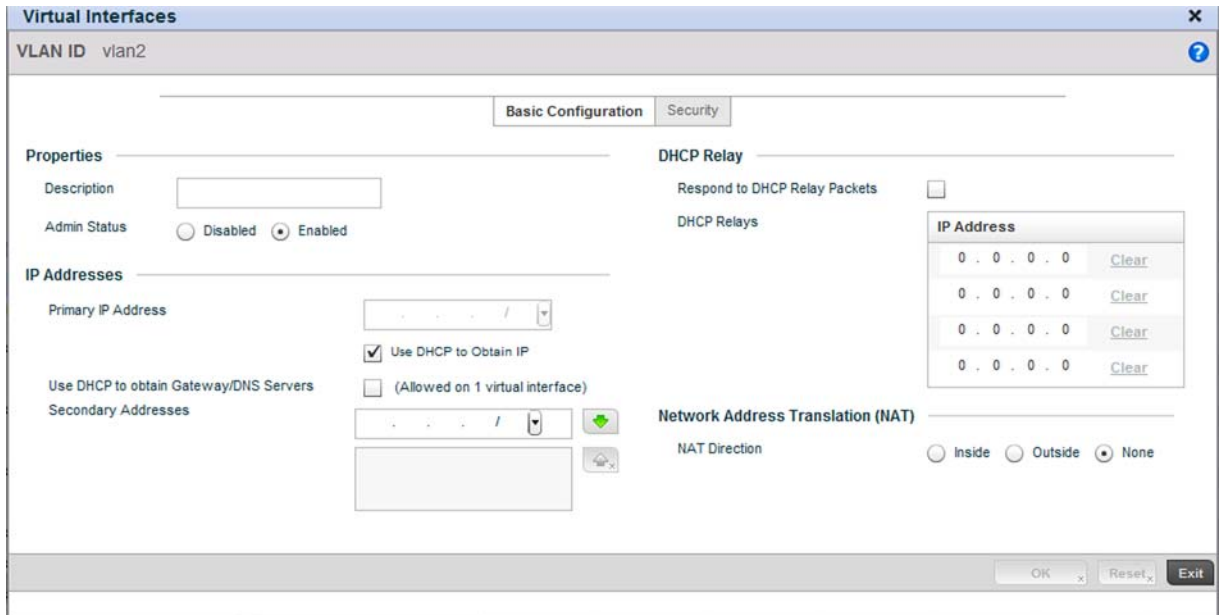


FIGURE 54 Virtual Interface screen

Define the **VLAN ID** for the AP7131 profile as VLAN 2.

Select the **Use DHCP to Obtain IP** option for setting the Primary IP Address.

Select **OK** to save the changes. Select **Exit** to close the screen, then click the **Commit** icon at the top right of the screen to apply the updates to the controller's running configuration.

Configure the Physical Interface

To configure the GE1 port:

Select and display the AP7131_UseCase1 profile by navigating to **Configuration > Profiles** and selecting the profile from amongst those displayed.

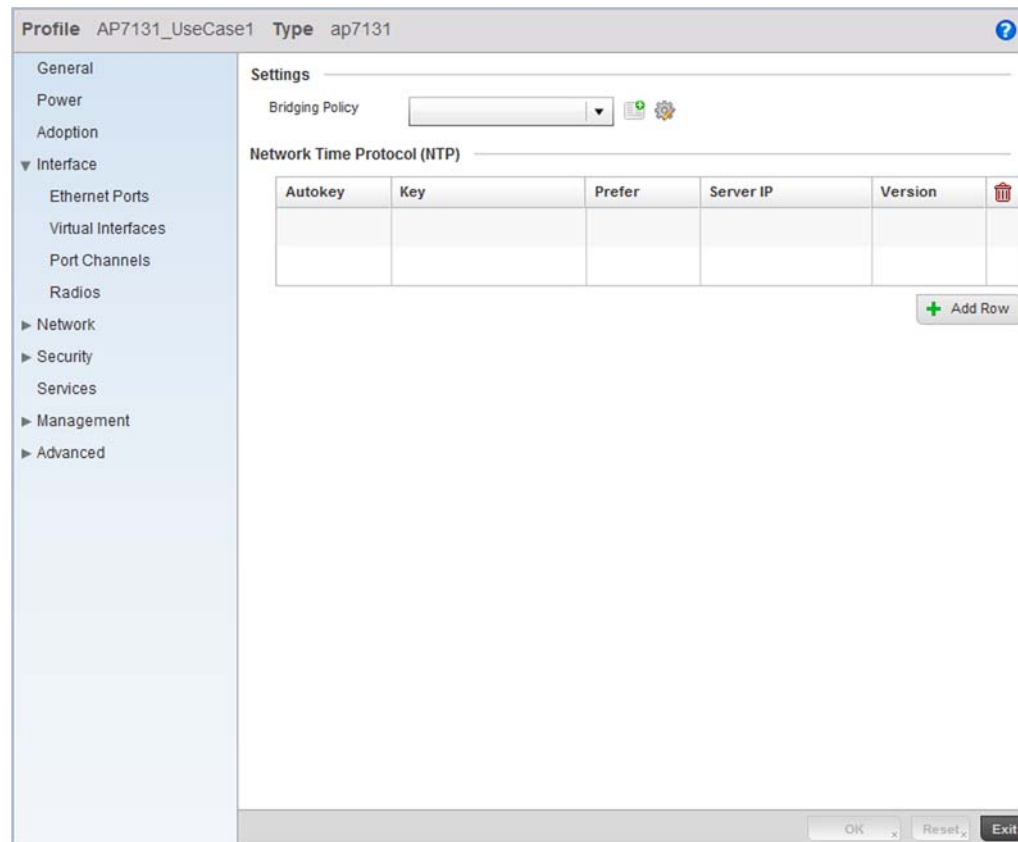


FIGURE 55 Brocade Mobility 71XX Access Point Profile Screen

Select the **Interface** menu option and expand it to display its submenu options.

Select **Ethernet Ports**.

The screenshot shows the 'Ethernet Ports' configuration window for interface 'ge1'. It has three tabs: 'Basic Configuration', 'Security', and 'Spanning Tree'. The 'Basic Configuration' tab is active. The 'Properties' section on the left contains: Description (text box with 'port1 description'), Admin Status (radio buttons for 'Disabled' and 'Enabled', with 'Disabled' selected), Speed (dropdown menu with '100'), Duplex (dropdown menu with 'Automatic'), and Allow Cisco Discover Protocol (checkbox checked). The 'Switching Mode' section on the right contains: Mode (radio buttons for 'Access' and 'Trunk', with 'Access' selected), Native VLAN (spinner control with '2' and range '(1 to 4,094)'), Tag Native VLAN (checkbox checked), and Allowed VLANs (text box with '1,3' and range '(2,4,7-12,...)'). The 'Port Channel Membership' section contains: Port Channel (checkbox unchecked and spinner control with '1' and range '(1 to 5)'). At the bottom right are buttons for 'OK', 'Reset', and 'Exit'.

FIGURE 56 Interface GE1 Configuration screen

Use the spinner control to define the **Native VLAN** as 2. No other values require configuration within the screen.

Select **OK** to save the changes. Select **Exit** to close the screen, then click the **Commit** icon at the top right of the screen to apply the updates to the controller's running configuration.

Repeat the steps in this section to configure the GE2 interface as well. Once completed, the screen appear as [Figure 57](#).

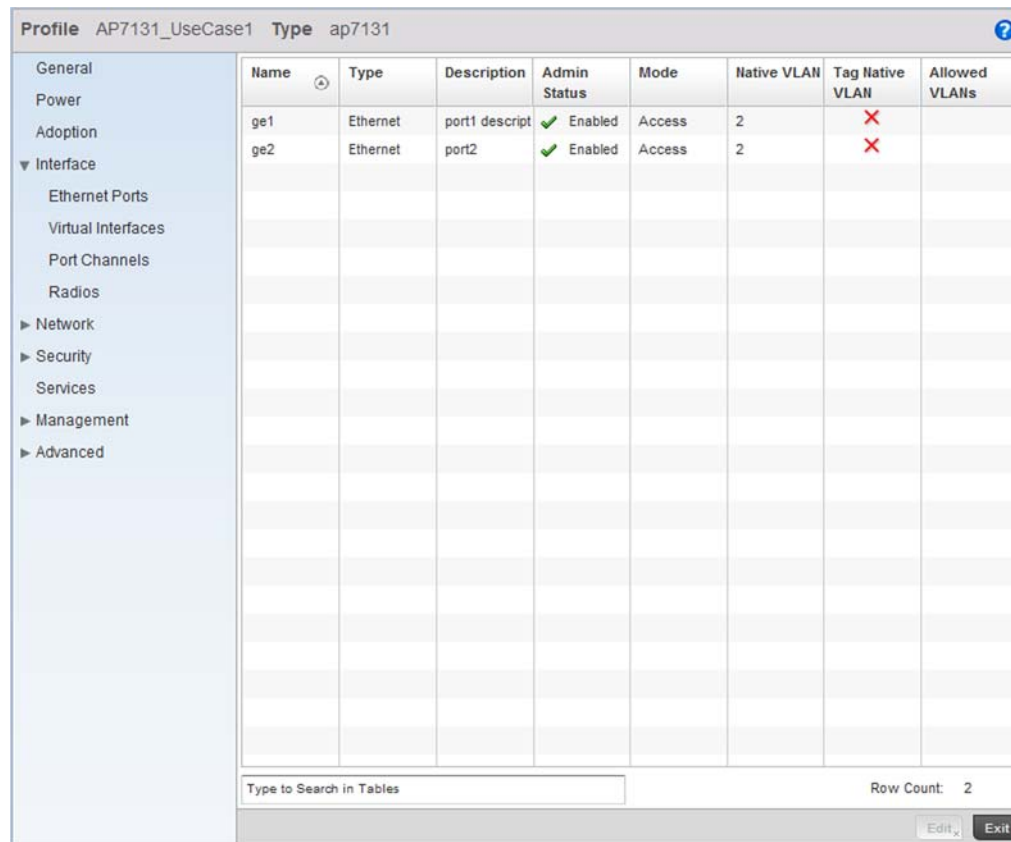


FIGURE 57 After configuring the Physical Interfaces on the Brocade Mobility 71XX Access Point Profile

Configure the Brocade Mobility 71XX Access Point's Radios

There are up to three radios within an AP7131 series Access Point (depending on the model purchased). However, the third AP7131 radio acts as a sensor and not available for WLAN support. Therefore, a maximum of two radios are available on an AP-7131 for WLAN support.

Each WLAN must be assigned an Access Point before wireless clients can connect to it. To configure an Brocade Mobility 71XX Access Point radios:

Select and display the AP7131_UseCase1 profile by navigating to **Configuration > Profiles** and selecting the profile from amongst those displayed.

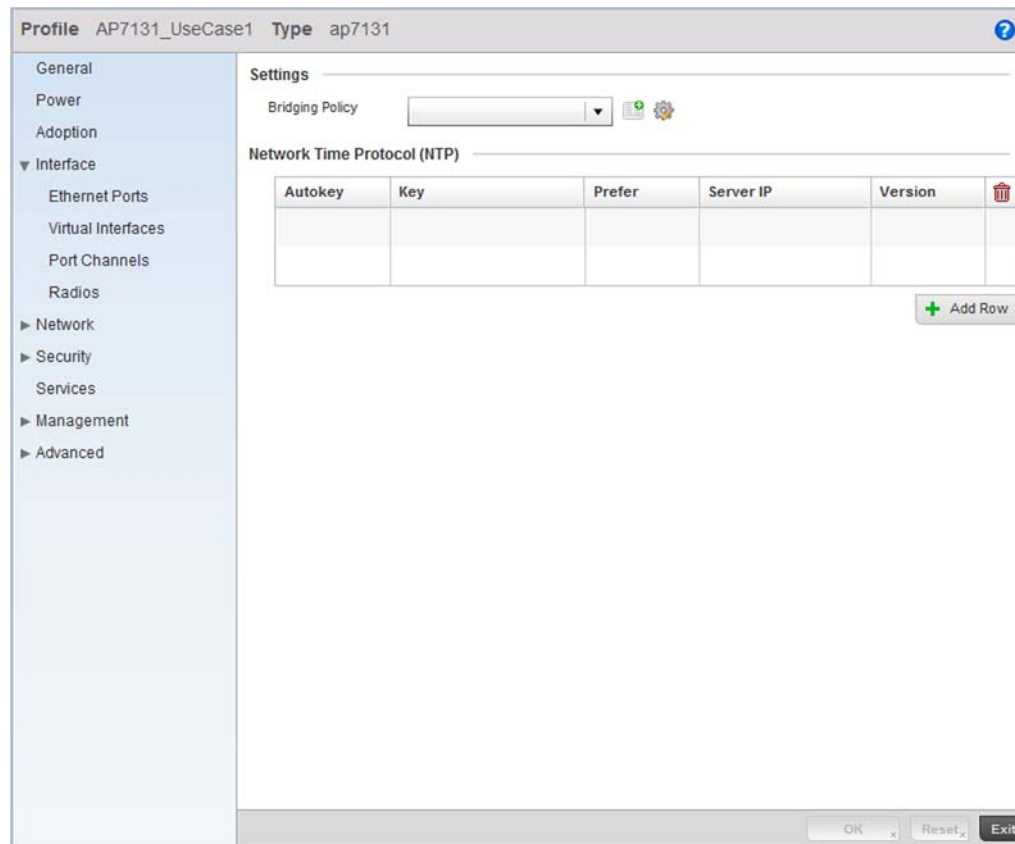


FIGURE 58 AP7131 Profile screen

Select the **Interface** menu item and expand it to display its submenu options.

Select **Radios**.

Profile AP7131_UseCase1 Type ap7131

General
Power
Adoption
▼ Interface
 Ethernet Ports
 Virtual Interfaces
 Port Channels
 Radios
► Network
► Security
 Services
► Management
► Advanced

Name	Type	Description	Admin Status	RF Mode	Channel	Transmit Power
radio1	Radio		✔ Enabled	2.4 GHz WLAN	smart	20
radio2	Radio		✔ Enabled	5 GHz WLAN	smart	20
radio3	Radio		✔ Enabled	Sensor	smart	20

Type to Search in Tables Row Count: 3

Edit Exit

FIGURE 59 Radios screen
Select **Radio1** and select the **Edit** button.

Radios

Name radio1

Radio Settings | WLAN Mapping | Advanced Settings

Properties

Description

Admin Status Disabled Enabled

Radio QoS Policy * default

Association ACL

Radio Settings

RF Mode 2.4GHz-wlan

Lock Radio Band

Channel

Transmit Power smart 20

Antenna Gain 0.00 (0.00 - 15.00 dBi)

Antenna Mode 1x3

Dynamic Chain Selection

Data Rates Custom::2,9,basic9,18

Radio Placement Indoor

WLAN Properties

Beacon Interval 50 (milliseconds)

DTIM Interval BSSID 1,7

RTS Threshold 2347 (1 to 2,347 bytes)

Short Preamble

Guard Interval Any

Probe Response Rate follow-probe-request

Probe Response Retry

Channel Scanning

Enable Off Channel Scan

OK Reset Exit

FIGURE 60 Radio 1 Configuration screen

Select the **WLAN Mapping** tab.

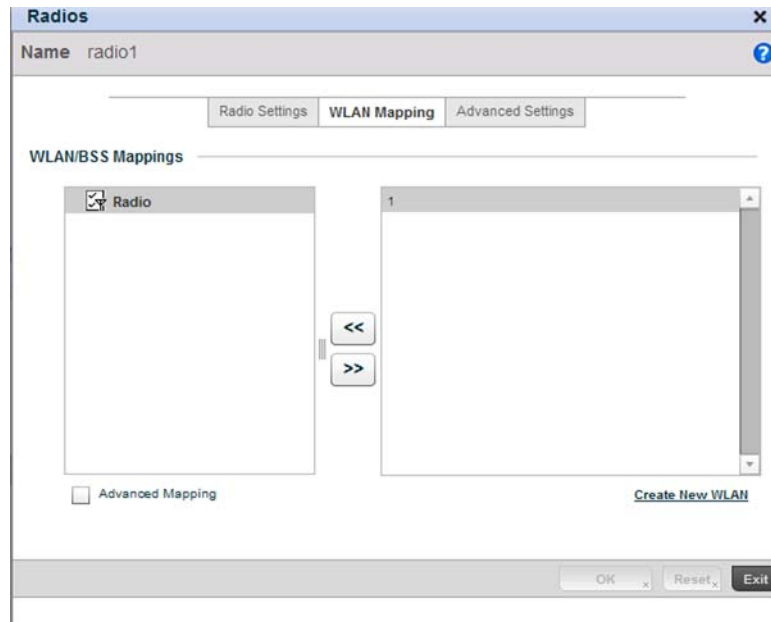


FIGURE 61 Radio 1 Configuration - WLAN Mapping screen

From the list on the right of the screen, select a WLAN to assign to this radio. Select the << button to assign the selected WLAN to Radio 1. The screen updates as displayed in [Figure 62](#).

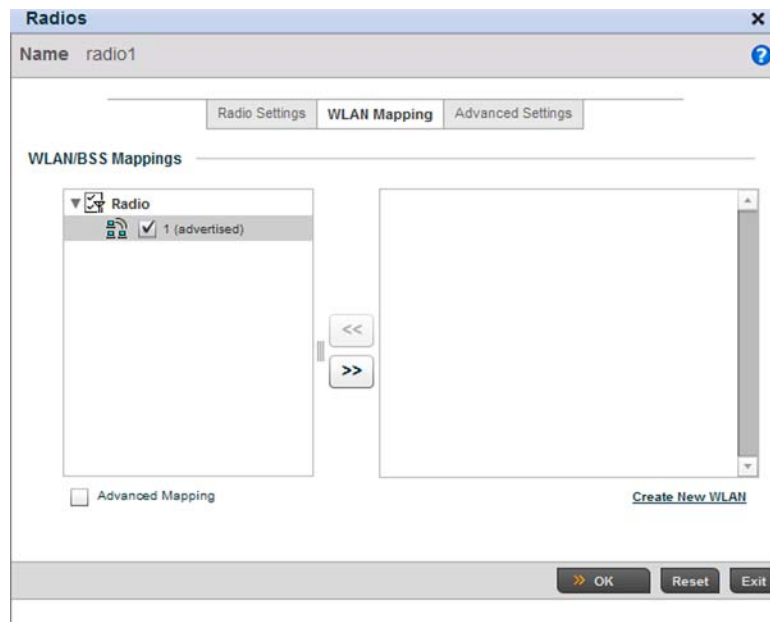


FIGURE 62 WLAN assigned to Radio1

Select **OK** to save the changes. Select **Exit** to exit the screen, then click the **Commit** icon at the top right of the screen to apply the updates to the controller's running configuration.

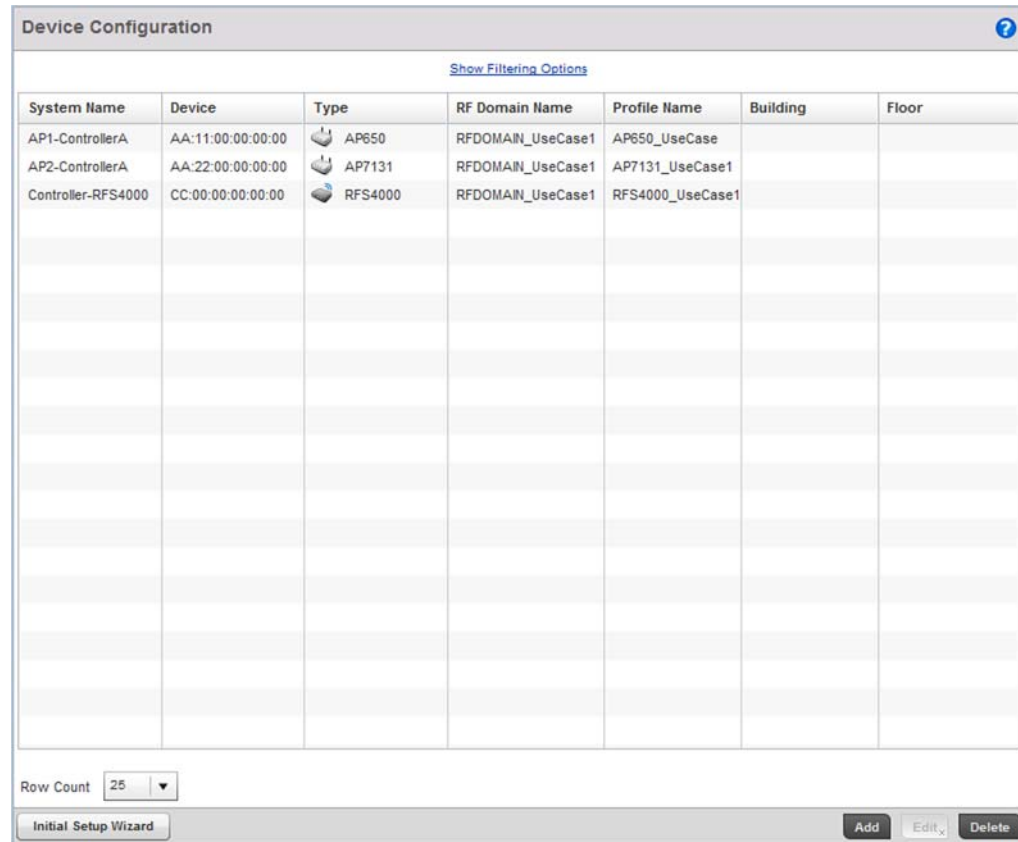
Configure the AP7131's Radio 2 by repeating steps 3 through 7.

Configure the Brocade Mobility 71XX Access Point to use the above profile

Before the Brocade Mobility 71XX Access Point can be utilized as a managed device with a WLAN, the profile must be applied to the access point.

To apply the profile to the Access Point:

Navigate to the Brocade Mobility 71XX Access Point profile by selecting **Configuration > Devices > Device Configuration**.



The screenshot shows a web interface titled "Device Configuration" with a help icon. Below the title is a "Show Filtering Options" link. The main content is a table with the following columns: System Name, Device, Type, RF Domain Name, Profile Name, Building, and Floor. The table contains three rows of data:

System Name	Device	Type	RF Domain Name	Profile Name	Building	Floor
AP1-ControllerA	AA:11:00:00:00:00	AP650	RFDOMAIN_UseCase1	AP650_UseCase		
AP2-ControllerA	AA:22:00:00:00:00	AP7131	RFDOMAIN_UseCase1	AP7131_UseCase1		
Controller-RFS4000	CC:00:00:00:00:00	RFS4000	RFDOMAIN_UseCase1	RFS4000_UseCase1		

At the bottom of the table, there is a "Row Count" dropdown menu set to "25". Below the table are three buttons: "Initial Setup Wizard", "Add", "Edit", and "Delete".

FIGURE 63 Device Configuration screen

Select the Brocade Mobility 71XX Access Point from amongst the list of devices displayed and select **Edit** button.

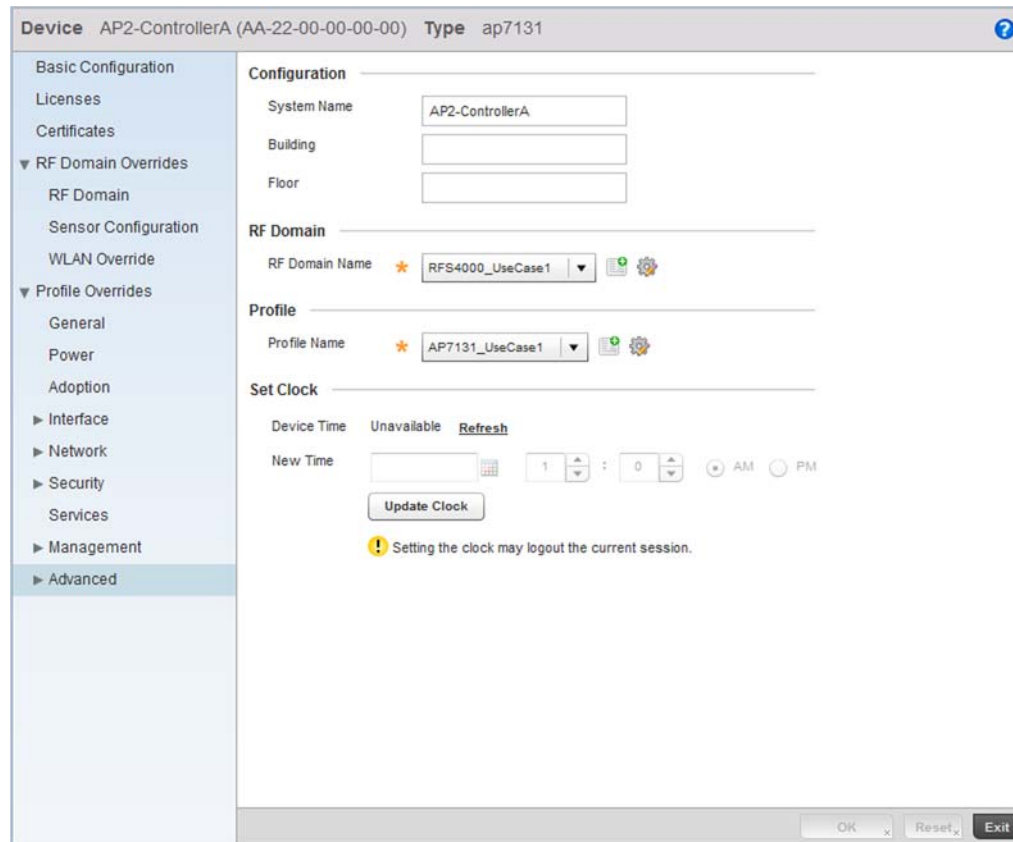


FIGURE 64 Brocade Mobility 71XX Access Point Profile Configuration screen

Select the **Profile Name** of Brocade Mobility 71XX Access Point_UseCase1 from the drop-down menu..

Select **OK** to save the changes. Select **Exit** to close this screen, then click the **Commit** icon at the top right of the screen to apply the updates to the controller's running configuration.

Creating a DHCP Server Policy

Using the Controller GUI to Configure the WLAN

The DHCP Server Policy sets the parameters required to run a DHCP server on the wireless controller and assign IP addresses automatically to device that associate. Configuring DHCP enables the reuse of a limited set of available IP addresses.

To create a DHCP server policy:

Select **Configuration > Services > DHCP Server Policy**.

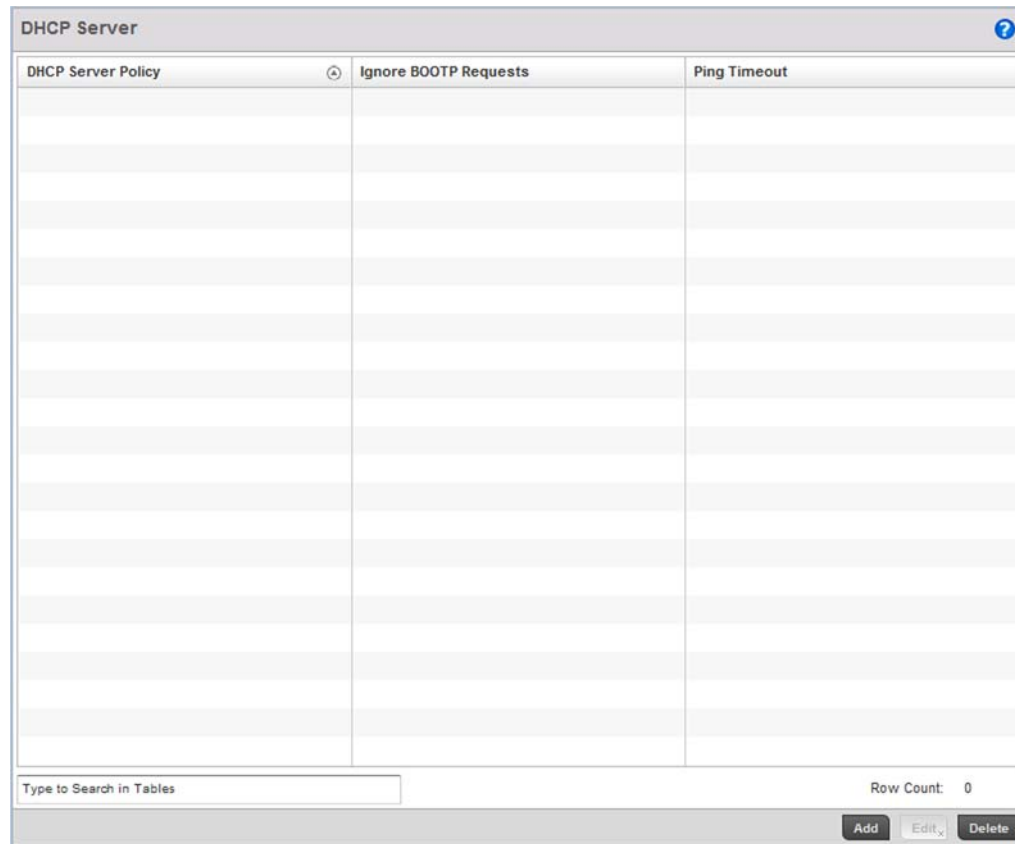


FIGURE 65 DHCP Server Policy screen

Select the **Add** button to create a new DHCP Policy.




FIGURE 66 DHCP Server Policy screen

Define a **DHCP Server Policy**. For the purposes of this use case scenario, use: DHCP_POLICY_UseCase1.

Select the **Continue** button to create this policy and enable the tabs required to configure its parameters.

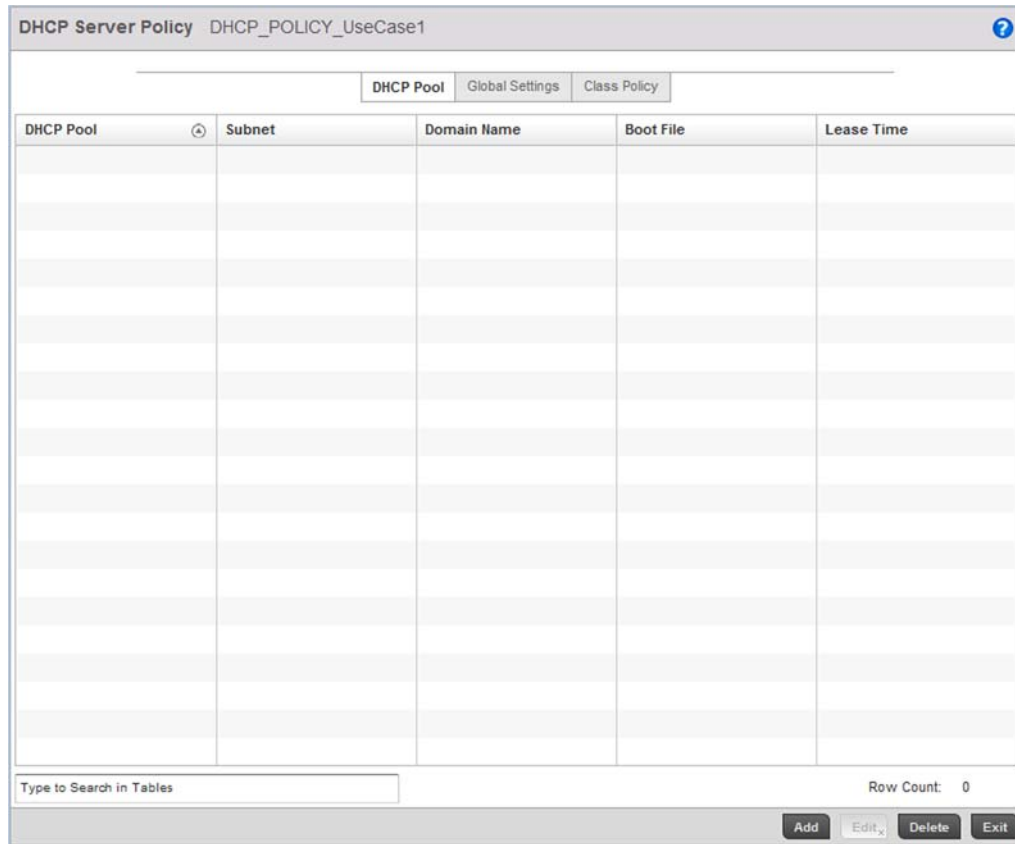


FIGURE 67 DHCP Pool screen

Select **Add** to create a new DHCP Pool. A DHCP Pool contains a range of IP addresses assigned to wireless clients and APs.

The screenshot shows the 'DHCP Pools' configuration window. The title bar indicates the pool name is 'DHCP_POOL_UseCase1_01'. The 'Basic Settings' tab is active. Under the 'General' section, the Subnet is set to '172.16.11.0 / 24'. The Lease Time is checked and set to '86400'. The 'IP Address Ranges' table contains one row with 'IP Start' 172.16.11.11 and 'IP End' 172.16.11.200. The 'Excluded IP Address Range' table is empty. At the bottom right, there are 'OK', 'Reset', and 'Exit' buttons.

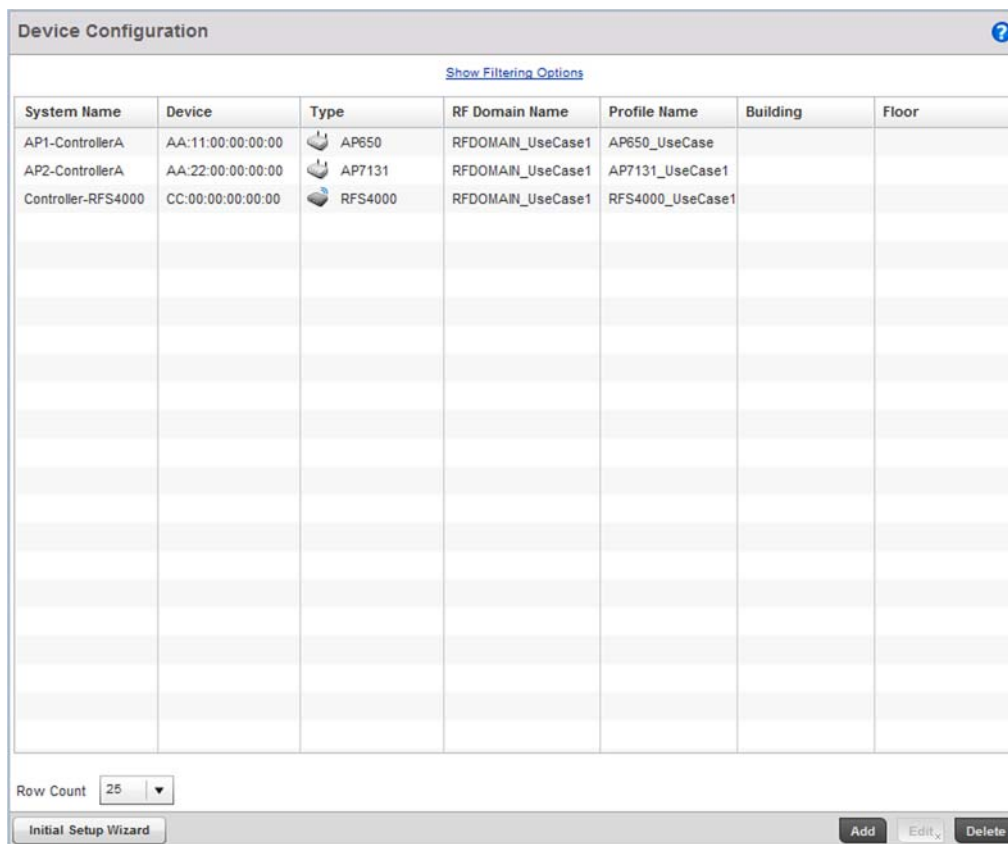
FIGURE 68 New DHCP Pool screen

Define the following parameters for the DHCP Pool configuration:

- DHCP Pool** Define the name of the DHCP Pool. For this scenario, use: DHCP_POOL_UseCase1_01
- Subnet** Assign the network on which this DHCP Server Policy is applied. For this scenario, use:
Value: 172.16.11.0/24
- IP Address Range** Provide the IP address range for this DHCP Pool. Select the **Add Row** button below this table to add a row. For this use case, use:
IP Start: 172.16.11.11
IP End: 172.16.11.200

Select **OK** to save the configuration. Select **Exit** to exit the screen, then click the **Commit** icon at the top right of the screen to apply the updates to the controller's running configuration.

Select the **Exit** button to return to the DHCP Server Policy screen. Once completed, the screen should appear as [Figure 69](#)



The screenshot shows a web interface titled "Device Configuration" with a help icon in the top right. Below the title is a link for "Show Filtering Options". The main content is a table with the following columns: System Name, Device, Type, RF Domain Name, Profile Name, Building, and Floor. The table contains three rows of data. Below the table, there is a "Row Count" dropdown menu set to 25, an "Initial Setup Wizard" button, and three buttons: "Add", "Edit", and "Delete".

System Name	Device	Type	RF Domain Name	Profile Name	Building	Floor
AP1-ControllerA	AA:11:00:00:00:00	AP650	RFDOMAIN_UseCase1	AP650_UseCase		
AP2-ControllerA	AA:22:00:00:00:00	AP7131	RFDOMAIN_UseCase1	AP7131_UseCase1		
Controller-RFS4000	CC:00:00:00:00:00	RFS4000	RFDOMAIN_UseCase1	RFS4000_UseCase1		

FIGURE 70 Device Configuration screen

Select the Brocade Mobility RFS4000 from the list and select **Edit**.

Device ControllerD-RFS4000 (DD-00-00-00-00-00) Type rfs4000

Basic Configuration

- Licenses
- Certificates
- RF Domain Overrides
 - RF Domain
 - Sensor Configuration
 - WLAN Override
- Profile Overrides
 - General
 - Cluster
 - Interface
 - Network
 - Security
 - Services
 - Management
 - Advanced

Configuration

System Name: ControllerD-RFS4000

Building: []

Floor: []

RF Domain

RF Domain Name: RFS4000_UseCase1

Profile

Profile Name: default-rfs4000

Set Clock

Device Time: Unavailable Refresh

New Time: [] : [] AM PM

Update Clock

Setting the clock may logout the current session.

OK Reset Exit

FIGURE 71 Brocade Mobility RFS4000 Device Context screen

From the menu on the left, select **Services**.

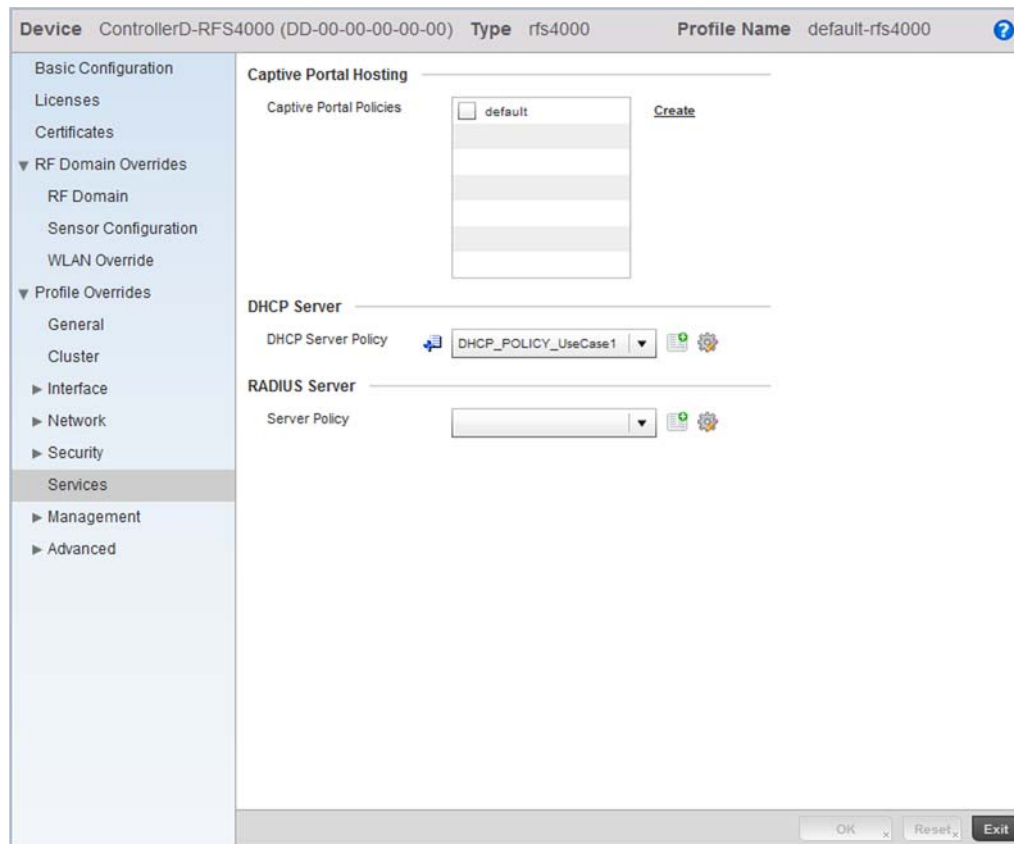


FIGURE 72 Brocade Mobility RFS4000 Device Context screen

From the **DHCP Server Policy** drop-down menu, select the name of the DHCP Server Policy. For this use case scenario, select DHCP_POLICY_UseCase1.

Select **OK** to save the changes. Select **Exit** to exit out of the Brocade Mobility RFS4000 wireless controller device context. Select **Commit** to save these changes to the configuration.

Completing and testing the configurations

Using the Controller GUI to Configure the WLAN

For a wireless client to successfully associate itself with the WLAN that you created, it must be configured. The following information must be used.

- **SSID:** WLAN_USECASE_01
- **Country:** Same as configured above in section *Creating a RF Domain* on page 3-26. In this example, the country code is set to US.
- **Mode:** Infrastructure

As the WLAN is set to beaconing, use the wireless client's wireless discovery client to discover the configured WLAN and associate to it.

Dashboard

In this chapter

- [Summary](#) 79
- [Network View](#) 88

The wireless controller dashboard enables wireless network administrators to review and troubleshoot the operation of the devices comprising the managed network. Additionally, the dashboard allows the review of the current network topology, the assessment of the network's component health and a diagnostic review of device performance.

By default, the Dashboard screen displays the System Dashboard, which is the top level in the device hierarchy.

Summary

The *Dashboard* displays device information organized by device association and inter-connectivity between the connected Access Points and wireless clients.

To review dashboard information, select **Dashboard** > **Summary**.

The Dashboard displays the **Health** tab by default.

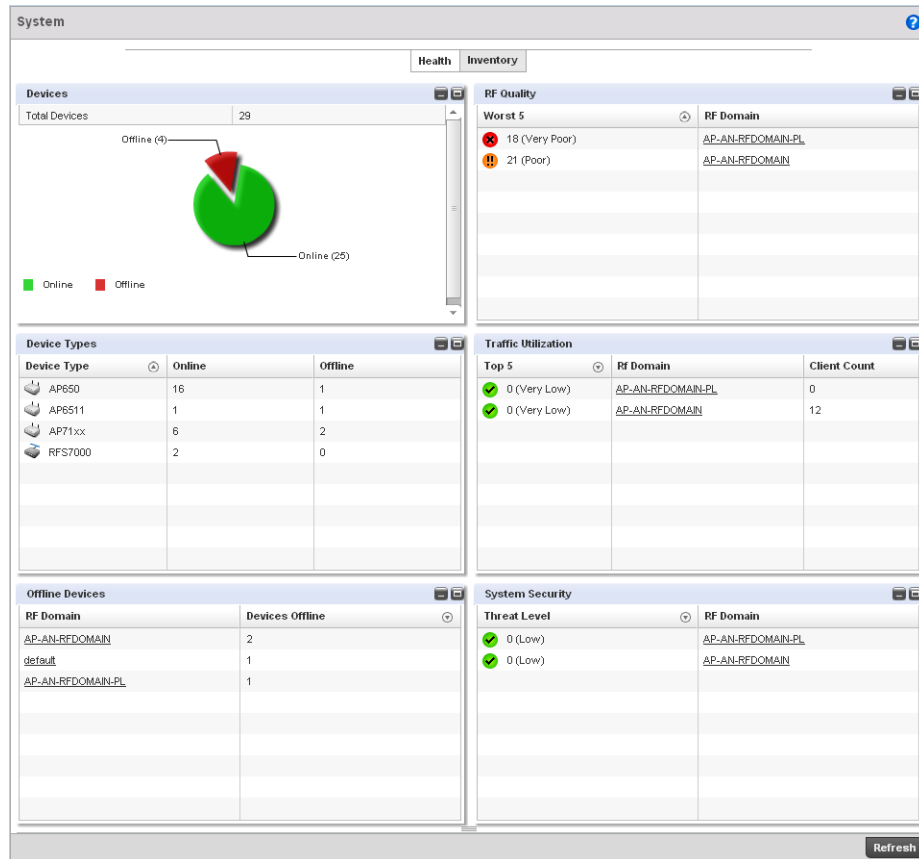


FIGURE 73 System Dashboard screen - Health tab

Device Listing

Summary

The device menu displays information as a hierarchical tree where each node is a RF Domain.

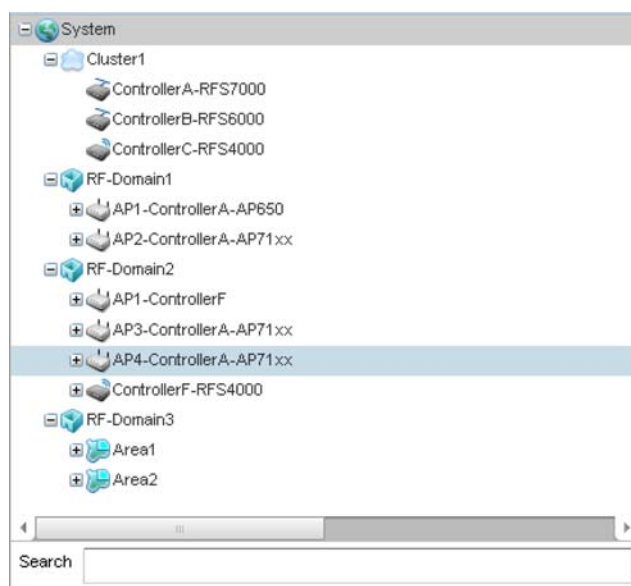


FIGURE 74 Dashboard Menu Tree

The **Search** text box, at the bottom, enables you to filter (search amongst) RF Domains. The **By** drop-down menu refines the search. You can further refine a search using the following:

- *Auto* – The search is automatically set to device type.
- *Name* – The search is performed for the device name specified in the **Search** text box.
- *WLAN* – The search is performed for the WLAN specified in the **Search** text box.
- *IP Address* – The search is performed for the IP Address specified in the **Search** text box.
- *MAC Address* – The search is performed for the MAC Address specified in the **Search** text box.

System Screen

The *System* screen displays the status of the managed network. The screen is partitioned into the following tabs:

- *Health* – The Health tab displays information about the state of the system being managed.
- *Inventory* – The Inventory tab displays information on the physical devices being managed by the system.

Health

Health

The **Health** tab displays device performance status for managed devices and includes their member RF Domains.

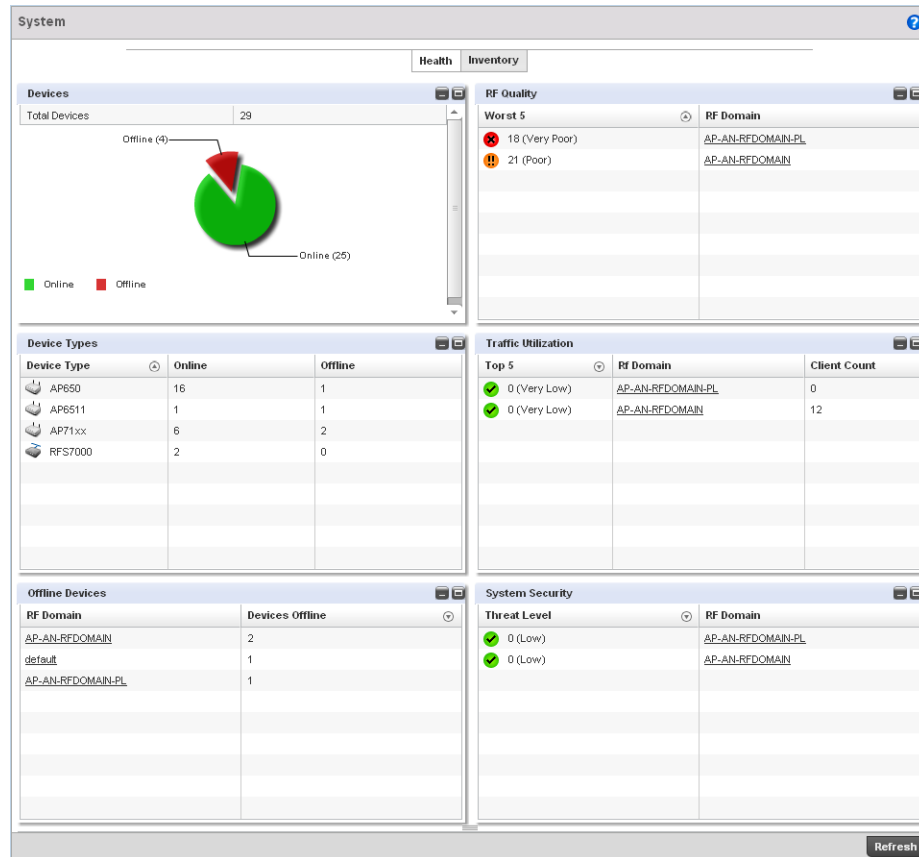


FIGURE 75 System Dashboard screen - Health tab

The Health screen is partitioned into the following fields:

- [Devices](#)
- [RF Quality Index](#)
- [Utilization](#)
- [Devices](#)
- [Wireless Clients](#)
- [Radios](#)
- [Client on Channels](#)

Devices

[Health](#)

The **Devices** field displays graphical status of the devices managed by this controller.

Device Details	
Hostname	ap6511-083571
Device MAC	5C-0E-8B-08-35-71
Type	AP6511
RF Domain Name	default
Version	5.0.1.0-044R
Uptime	8 days, 11 hours 57 minutes
CPU	MIPS 24Kc V7.4
RAM	55260 kB
System Clock	Nov 19 22:31:28 PST 2010

FIGURE 76 System Dashboard screen - Health tab - Device Health field

The Devices field displays the total device count managed by this wireless controller and their status (online vs. offline) in pie chart format. Use this information to determine whether the number of offline devices requires troubleshooting to improve the performance of the controller managed network.

RF Quality Index

Health

The RF Quality Index displays RF quality per RF Domain. It's a measure of the overall effectiveness of the RF environment displayed in percentage. It's a function of the connect rate in both directions, retry rate and error rate.

Radio RF Quality Index		
RF Quality Index	Radio Id	Radio Type
✓ 100 (Good)	5C-0E-8B-08-35-71:R1	2.4 GHz WLAN

FIGURE 77 System Dashboard screen - Health tab - RF Quality Index field

The RF Quality field displays an average quality index supporting all the RF Domains on the wireless controller. The table lists the bottom five (5) RF quality values for RF Domains supported on the wireless controller.

The quality is measured as:

- 0-20 – Very poor quality
- 20-40 – Poor quality
- 40-60 – Average quality
- 60-100 – Good quality

Select an RF Domain to view its performance statistics.

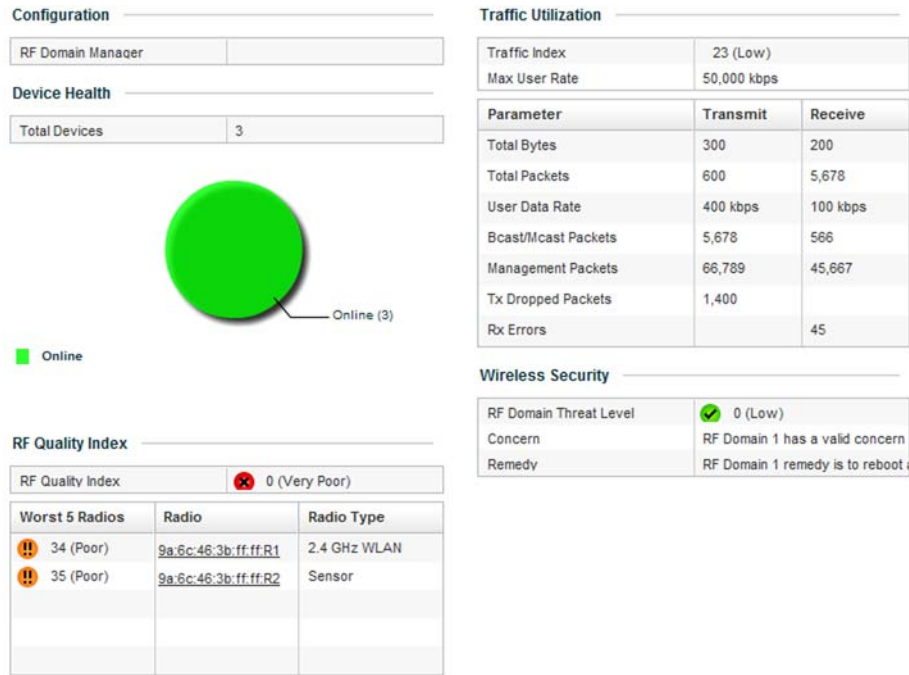


FIGURE 78 RF Quality Index - RF Domains

Select a RF Domain to review poorly performing radios.

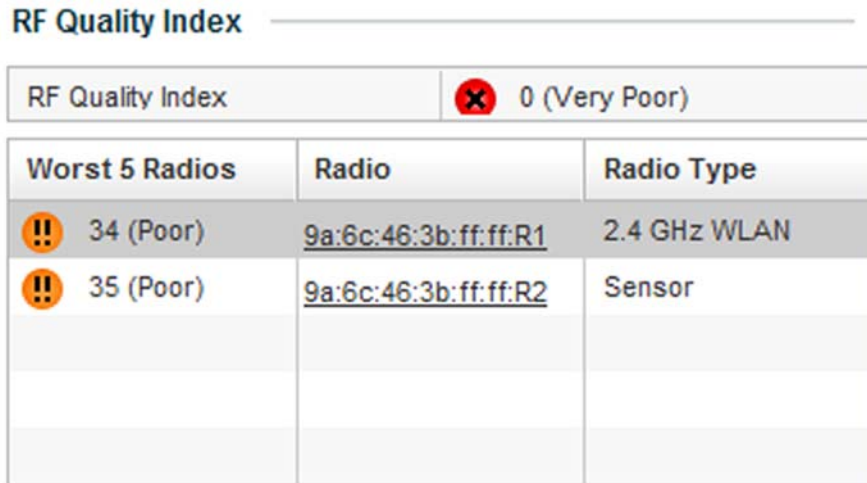


FIGURE 79 RF Quality Index - Worst Performing Radios

The following screen displays.

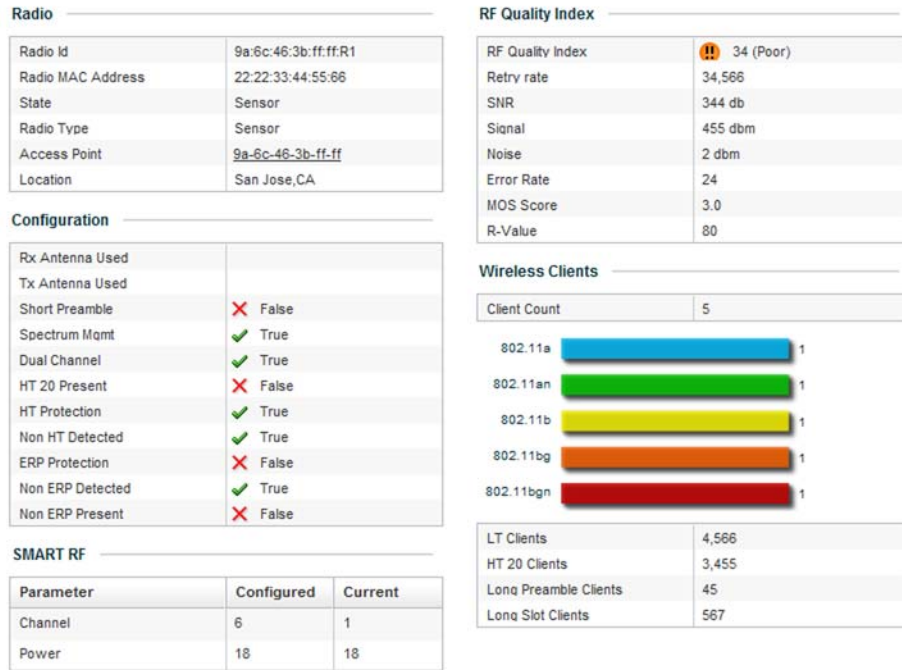


FIGURE 80 RF Quality Index - Radio Statistics

Use this diagnostic information to determine what measures can be taken to improve radio performance in respect to wireless client load and the radio bands supported.

For information on RF Domains, and how to create one for use with the managed network, see [Managing RF Domains on page 8-408](#).

Utilization

Health

The Utilization field displays RF medium efficiency. Traffic utilization is the percentage of current throughput relative to the maximum possible throughput for a managed RF Domain.

The Utilization field displays a list of up to five RF Domains in relation to the number of associated wireless clients. It also displays a table of the packets types transmitted.

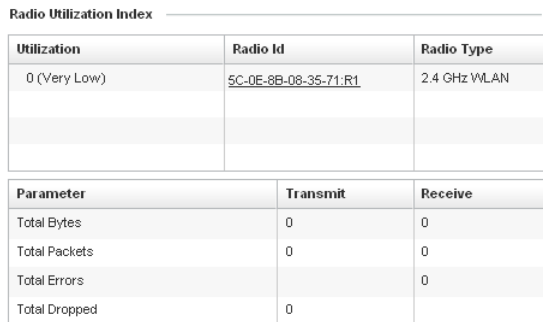


FIGURE 81 System Dashboard screen - Health tab - Utilization field

Inventory

System Screen

The System screen's *Inventory* tab displays information on devices managed by this system. The screen provides a complete overview of the number and state of devices managed by the system. Information is displayed in easy to read tables and graphs. This screen also provides links for more detailed information.

To navigate to this screen, select **Dashboard > Dashboard > RF Domain > Network**.

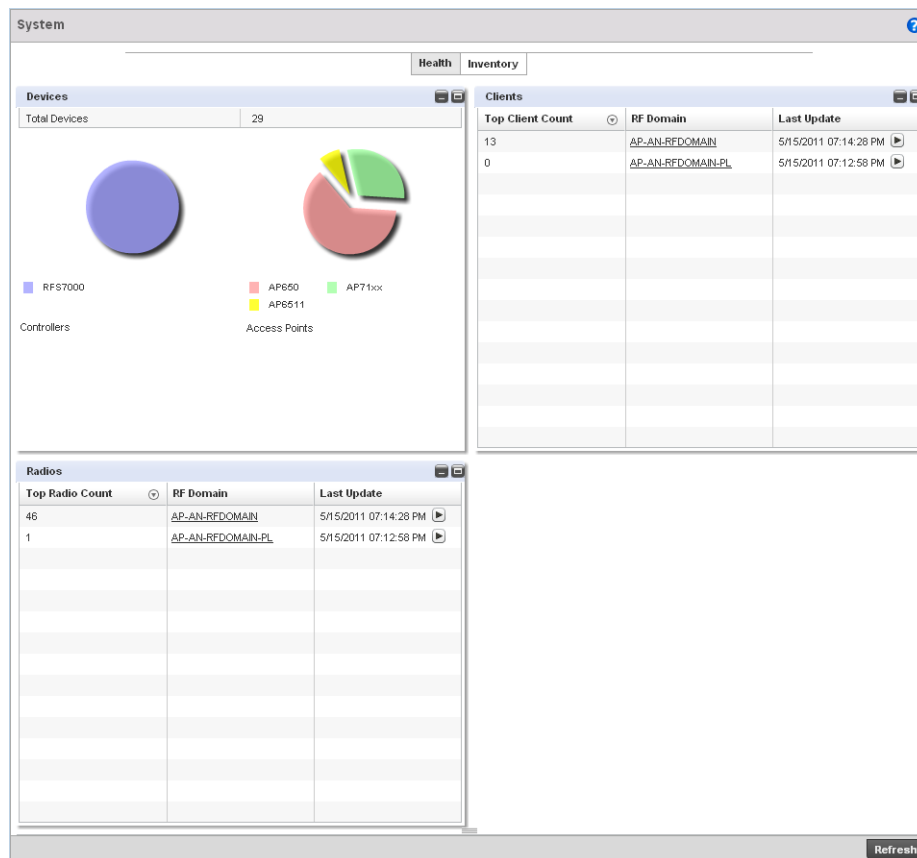


FIGURE 82 System screen - Inventory tab

The information within the Inventory tab is partitioned into the following fields:

- [Devices](#)
- [Wireless Clients](#)
- [Radios](#)
- [Client on Channels](#)

Devices

Inventory

The Devices field displays a ratio of peer controllers and managed Access Points. The information is displayed in pie chart format.

Radio Types

Total Radios	1
--------------	---

2.4 GHz WLAN  1

FIGURE 83 System screen - Inventory tab - Device Types field

The Device Type field displays a numerical representation of the different controllers models and connected Access Points in the current system. Does this device distribution adequately support the number and types of Access Points and their client load.

Clients

Inventory

The Clients field displays information about the wireless clients managed by the controller's connected Access Point radios.

Wireless Clients

Total Wireless Clients	0	
Top Client Count	Radio	Radio Type
0	5C-0E-8B-08-35-59.R1	2.4 GHz WLAN

FIGURE 84 System screen - Inventory tab - Wireless Clients field

Information in the Wireless Clients field displays in two tables. The first lists the total number of wireless clients managed by this system. The second lists the top five RF Domains in respect to the number of connected clients.

Each RF Domain can be selected and analyzed in respect to its performance. For information on RF Domains, and how to create one for use with the managed network, see [Managing RF Domains on page 8-408](#).

Radios

Inventory

The Radios field displays information about the different radios managed by this system.

Wireless Clients

Total Wireless Clients	0	
Top Client Count	Radio	Radio Type
0	5C-0E-8B-08-35-59.R1	2.4 GHz WLAN

FIGURE 85 System screen - Inventory tab - Radios field

Information in the Radio area is presented in two tables. The first lists the total number of Radios managed by this system, the second lists the top five RF Domains in terms of the number of available radios.

Client on Channels

Inventory

The Client on Channels field displays bar-graphs of wireless clients classified by channel and radio band.




FIGURE 86 System screen - Inventory tab - Clients on Channels field

Wireless Clients are displayed as either operating in the 5 GHz or 2.4 GHz channel. Information is further classified by radio band. For the 5 GHz channel, information is classified by either the 802.11a or 802.11an bands. For the 2.4 GHz channel, information is classified by either the 802.11b, 802.11bg or 802.11bgn bands. Does this client distributions adequately support the requirements of the radio coverage area?

Network View

The wireless controller's Network View functionality displays device association connectivity amongst the wireless controller, access point and wireless clients. This association is represented by a number of different graphs.

To review the wireless controller's Network Topology, select **Dashboard > Overview > Network**.

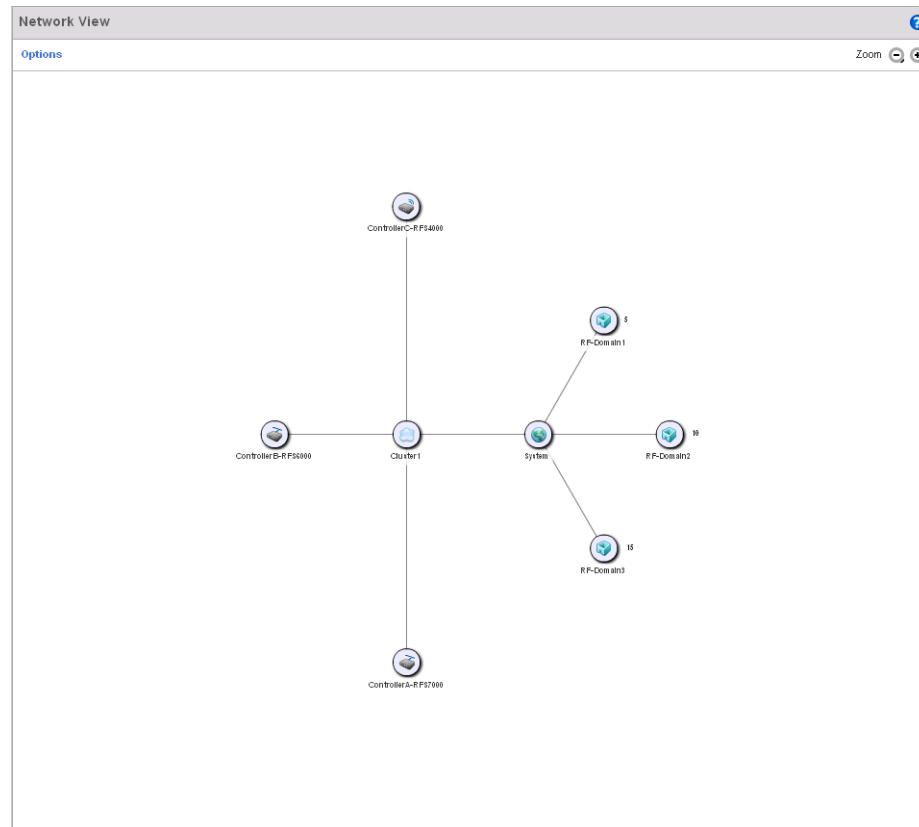


FIGURE 87 Network View Topology

The screen displays icons for the different views available to the system. Apart from device specific icons, the following three icons are available:

- *default* – Displays information about the default RF Domain.
- *system* – Displays information about the current system.
- *cluster* – Displays information about clusters managed by this system.

Use these icons to navigate quickly within top level groupings.

The middle field displays a *Network View*, or graphical representation of the network. This field changes to display a graphical network map.

Select the **Settings** link (the blue link near the top of the screen) to define how devices are displayed within the Network View.

Settings		
Display		
	Show Item	Show Label
Clients	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Access Points	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Ok Cancel

FIGURE 88 Network View - Settings field

Select either or both of the **Access Point** and **Client** options to display them in the Network View. Similarly, select the **Show Label** option to display hardware MAC address as an appended label. Select **OK** to save the updates.

Device Configuration

In this chapter

• Basic Configuration	92
• Basic Device Configuration	94
• License Configuration	96
• Assigning Certificates	98
• RF Domain Overrides	115
• Profile Overrides	121
• Auto Provisioning Policies	216
• Critical Resource Policy	222

Devices managed by the controller can either be assigned unique configurations or have existing RF Domain or Profile configurations modified (overridden) to support a requirement that dictates a device's configuration be customized from the configuration shared by its peer devices.

When a device is initially managed by the controller, it requires several basic configuration parameters be set (system name, deployment location etc.). Additionally, the number of permitted device licenses (purchased directly from Brocade Mobility) needs to be accessed to determine whether a new *Access Point* (AP) or *Adaptive Access Point* (AAP) can be adopted.

Refer to the following to set a device's basic configuration, license and certificate usage:

- [Basic Configuration](#)
- [Basic Device Configuration](#)
- [License Configuration](#)
- [Assigning Certificates](#)

RF Domains allow administrators to assign configuration data to multiple devices deployed in a common coverage area (floor, building or site). In such instances, there's many configuration attributes these devices share as their general client support roles are quite similar. However, device configurations may need periodic refinement (overrides) from their original RF Domain administered design. For more information, see *RF Domain Overrides on page 5-115*.

Profiles enable administrators to assign a common set of configuration parameters and policies to controllers and Access Points. Profiles can be used to assign shared or unique network, wireless and security parameters to wireless controllers and Access Points across a large, multi segment, site. The configuration parameters within a profile are based on the hardware model the profile was created to support. The controller supports both default and user defined profiles implementing new features or updating existing parameters to groups of wireless controllers or Access Points.

However, device profile configurations may need periodic refinement from their original administered configuration. Consequently, a device profile could be applied an override from the configuration shared amongst numerous peer devices deployed within a particular site. For more information, see *Profile Overrides on page 5-121*.

Adoption is the process an Access Point uses to discover controllers available in the network, pick the most desirable controller, establish an association, obtain its configuration and consider itself provisioned.

At adoption, an Access Point solicits and receives multiple adoption responses from controllers available on the network. Modify existing adoption policies or create a new one as needed to meet the adoption requirements of a device and its assigned controller profile. For more information, see *Auto Provisioning Policies on page 5-216*.

Lastly, use **Configuration > Devices** to define and manage a critical resource policy. A critical resource policy defines a list of device IP addresses on the network (gateways, routers etc.). The support of these IP address is interpreted as critical to the health of the managed network. These devices addresses are pinged regularly by the controller. If there's a connectivity issue, an event is generated stating a critical resource is unavailable. For more information, see *Critical Resource Policy on page 5-222*.

Basic Configuration

To assign a Basic Configuration:

1. Select the **Configuration** tab from the Web UI.
2. Select **Devices** from the Configuration tab.
3. The Device Configuration screen displays a list of managed devices or other controllers.

Device Configuration								
System Name	Device	Type	RF Domain Name	Profile Name	Area	Floor	Overrides	
AP1-ControllerA-AP650	AA-11-00-00-00-00	AP650	RF-Domain1	default-ap650	Not Set	Not Set		
AP1-ControllerB-AP650	BB-11-00-00-00-00	AP650	RF-Domain3	default-ap650	Not Set	Not Set	Clear	
AP2-ControllerA-AP710	AA-22-00-00-00-00	AP71XX	RF-Domain1	default-ap71xx	Not Set	Not Set	Clear	
AP2-ControllerB-AP650	BB-22-00-00-00-00	AP650	RF-Domain3	default-ap650	Not Set	Not Set	Clear	
AP3-ControllerA-AP710	AA-33-00-00-00-00	AP71XX	RF-Domain2	default-ap71xx	Not Set	Not Set	Clear	
AP3-ControllerB-AP650	BB-33-00-00-00-00	AP650	RF-Domain3	default-ap650	Not Set	Not Set	Clear	
AP4-ControllerA-AP710	AA-44-00-00-00-00	AP71XX	RF-Domain2	default-ap71xx	Not Set	Not Set		
AP4-ControllerB-AP651	BB-44-00-00-00-00	AP6511	RF-Domain3	default-ap6511	Not Set	Not Set		
AP5-ControllerB-AP621	BB-55-00-00-00-00	AP621	RF-Domain3	default-ap621	Not Set	Not Set	Clear	
AP6-ControllerB-AP652	BB-66-00-00-00-00	AP6521	RF-Domain3	default-ap6521	Not Set	Not Set	Clear	
AP7-ControllerB-AP653	BB-77-00-00-00-00	AP6532	RF-Domain3	default-ap6532	Not Set	Not Set	Clear	
ControllerA-RFS7000	AA-00-00-00-00-00	RFS7000	RF-Domain1	default-rfs7000	Not Set	Not Set		
ControllerB-RFS6000	BB-00-00-00-00-00	RFS6000	RF-Domain1	RFS6000	Not Set	Not Set	Clear	
ControllerC-RFS4000	CC-00-00-00-00-00	RFS4000	RF-Domain3	default-rfs4000	Not Set	Not Set	Clear	
ControllerD-RFS4000	DD-00-00-00-00-00	RFS4000	RF-Domain3	default-rfs4000	Not Set	Not Set		

Type to search in tables Row Count: 15

Initial Setup Wizard Add Edit Delete

FIGURE 89 Device Configuration screen

Refer to the following device settings to determine whether a configuration update or RF Domain or Profile change is warranted:

System Name	Displays the name assigned to the device when the basic configuration was defined. This is also the device name that appears within the RF Domain or Profile the device supports.
Device	Displays the device's factory assigned MAC address used as hardware identifier. The MAC address cannot be revised with the device's configuration.
Type	Displays the Brocade Mobility device model for the listed Access Point or wireless controller.
RF Domain Name	Lists RF Domain memberships for each listed device. Devices can either belong to a default RF Domain based on model type, or be assigned a unique RF Domain supporting a specific configuration customized to that device model.
Profile Name	Lists the profile each listed device is currently a member of. Devices can either belong to a default profile based on model type, or be assigned a unique profile supporting a specific configuration customized to that model.

Area	List the physical area where the controller or access point is deployed. This can be a building, region, campus or other area that describes the deployment location.
Floor	List the building Floor name representative of the location within the area or building the controller or Access Point was physically deployed. Assigning a building Floor name is helpful when grouping devices in RF Domains and Profiles, as devices on the same physical building floor may need to share specific configuration parameters in respect to radio transmission and interference requirements specific to that location.
Overrides	The Overrides column contains an option to clear all profile overrides for any devices that contain overrides. To clear an override, select the clear button to the right of the device.

4. Select **Add** to create a new device. Select **Edit** to modify an existing device and select **Delete** to remove an existing device.

Basic Device Configuration

Setting a device's Basic Configuration is required to assign a device name, deployment location, and system time. Similarly, the Basic Configuration screen is where Profile and RF Domain assignments are made. RF Domains allow administrators to assign configuration data to multiple devices deployed in a common coverage area, such as in a floor, building or site. Each RF Domain contains policies that can determine a Smart RF or WIPS configuration.

Profiles enable administrators to assign a common set of configuration parameters and policies to controllers and Access Points. Profiles can be used to assign common or *unique* network, wireless and security parameters to wireless controllers and Access Points across a large, multi segment, site. The configuration parameters within a profile are based on the hardware model the profile was created to support. The controller supports both default and user defined profiles implementing new features or updating existing parameters to groups of wireless controllers or Access Points. The central benefit of a profile is its ability to update devices collectively without having to modify individual device configurations one at a time.

NOTE

Once devices have been assigned membership in either a profile or RF Domain, an administrator must be careful not to assign the device a configuration update that removes it from membership from the RF Domain or profile. A RF Domain or profile configuration must be re-applied to a device once its configuration has been modified in a manner that differentiates it from the configuration shared by the devices comprising the RF Domain or profile.

To assign a device a Basic Configuration:

1. Select the **Configuration** tab from the Web UI.
2. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

3. Select a target device (by double-clicking it) from amongst those displayed.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

4. The **Basic Configuration** screen displays by default.

FIGURE 90 Basic Configuration screen

- Set the following configuration settings for the target device:

System Name Provide the selected device a system name up to 64 characters. This is the device name that appears within the RF Domain or Profile the device supports.

Area Assign the device an Area name representative of the location the controller or Access Point was physically deployed. The name cannot exceed 64 characters. Assigning an area name is helpful when grouping devices in RF Domains and profiles, as devices in the same physical deployment location may need to share specific configuration parameters in respect to radio transmission and interference requirements specific to that location.

Floor Assign the target a device a building Floor name representative of the location the Access Point was physically deployed. The name cannot exceed 64 characters. Assigning a building Floor name is helpful when grouping devices in Profiles, as devices on the same physical building floor may need to share specific configuration parameters in respect to radio transmission and interference requirements specific to that location.

- Use the **RF Domain** drop-down menu to select an existing RF Domain for device membership. If a RF Domain configuration does not exist suiting the deployment requirements of the target device, select the **Create** icon to create a new RF Domain configuration, or select the **Edit** icon to modify the configuration of a selected RF Domain. For more information, see [Managing RF Domains on page 8-408](#).
- Use the **Profile** drop-down menu to select an existing RF Domain for device membership.

If a profile configuration does not exist suiting the deployment requirements of the target device, select the **Create** icon to create a new profile configuration, or select the **Edit** icon to modify the configuration of a selected profile. For more information, see [General Profile Configuration on page 7-316](#).

8. If necessary, click the **Clear Overrides** button to remove all existing overrides from the device.
9. Refer to the **Set Clock** parameter to update the system time of the target device.
10. Refer to the **Device Time** parameter to assess the device's current time, or whether the device time is unavailable. Select **Refresh** as required to update the device's reported system time.
11. Use the **New Time** parameter to set the calendar day, hour and minute for the target device. Use the *AM* and *PM* radio buttons to refine whether the updated time is for the morning or afternoon/evening.
12. When completed, select **Update Clock** to commit the updated time to the target device.
13. Select **OK** to save the changes made to the device's Basic Configuration. Selecting **Reset** reverts the screen to its last saved configuration.

License Configuration

Licenses are purchased directly from Brocade Mobility for the number of permissible *Access Point* (AP) and *Adaptive Access Point* (AAP) adoptions per controller or managed cluster.

NOTE

The Licenses screen is only available to wireless controllers capable of sustaining device connections, and thus require license support to provide the terms for the maximum number of allowed device connections. The License screen is not available for Access Points.

The Licenses screen also contains a facility where new licenses can be applied to increase the number of device adoptions permitted, or to allow the use of the advanced security or advanced WIPS features.

To configure a device's a license configuration:

1. Select the **Configuration** tab from the Web UI.
2. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

3. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

4. Select **Licenses** from the Device menu options.

Device Serial Number	
Serial Number	123456789012

Licenses	
AAP Adoptions	Unavailable
AAP Licenses	Unavailable
AP Adoptions	Unavailable
AP Licenses	Unavailable
Maximum APs	Unavailable

Cluster Licenses	
Cluster AAP Adoptions	Unavailable
Cluster AAP Licenses	Unavailable
Cluster AP Adoptions	Unavailable
Cluster AP Licenses	Unavailable
Cluster Maximum APs	Unavailable

Apply Licenses	
AP Licenses	<input type="text"/>
Adaptive AP Licenses	<input type="text"/>
Advanced Security	<input type="text"/>
Advanced WIPS Licenses	<input type="text"/>

OK Reset Exit

FIGURE 91 Device Licenses screen

The License screen displays a **Device Serial Number** of the controller used for generating the license key.

- Review the **Licenses** table, to assess the specific number of AP and AAP adoptions permitted, as dictated by the terms of the current license.

AAP Adoptions

Lists the total number of AAP adoptions currently made by the target controller. If the installed license count is 10 AAPs and the number of AAP adoptions is 5, 5 additional AAPs can still be adopted under the terms of the current license. The total number of AAPs adoptions varies by controller platform, as well as the terms of the license.

AAP Licenses

Lists the number of AAPs available for adoption by the controller under the restrictions of the current license. This number applies to dependent mode adaptive APs only, and not independent mode APs.

AP Adoptions

Lists the total number of AP adoptions currently made by the target controller. If the installed license count is 20 APs and the number of AP adoptions is 10, 10 additional APs can still be adopted under the terms of the current license. The total number of APs adoptions varies by controller platform, as well as the terms of the license.

AP Licenses

Lists the number of APs available for adoption by the controller under the restrictions of the current license. This number applies to independent mode APs only, and not dependent mode AAPs.

Maximum APs

Lists the maximum number of APs that can be supported by the listed controller under the terms of the license.

- Review the **Cluster Licenses** table, to assess the specific number of AP and AAP adoptions per controller cluster, as dictated by the terms of the current license.

Cluster AAP Adoptions	Lists the total number of AAP adoptions currently made by the target controller's cluster membership (includes all controller members). If the installed license count is 100 AAPs and the number of AAP adoptions is 50, 50 additional AAPs can still be adopted under the terms of the current AAP licenses, pooled by the cluster members.
Cluster AAP Licenses	Lists the number of AAPs available for adoption by the cluster member controllers under the restrictions of the licenses combined amongst the cluster members.
Cluster AP Adoptions	Lists the total number of AP adoptions currently made by the target controller's cluster membership (includes all controller members). If the installed license count is 100 APs and the number of AP adoptions is 40, 60 additional APs can still be adopted under the terms of the current AP licenses pooled by the cluster members.
Cluster AP Licenses	Lists the number of APs available for adoption by the cluster member controllers under the restrictions of the licenses accumulated amongst the cluster members.
Cluster AP Maximum APs	Lists the maximum number of cluster AP adoptions that can be supported by the listed controller or Access Point Controller under the terms of the license.

- Refer to the **Apply Licenses** field to apply licenses to APs and AAPs counts, as well as the provisioning of advanced security and advanced WIPS features:

AP Licenses	Enter the Brocade Mobility provided license key required to adopt a specified number of APs to the controller. The available number of AP licenses varies by controller platform.
Adaptive AP Licenses	Enter the Brocade Mobility provided license key required to install a specified number of AAPs to the controller. The available number of AAP licenses varies by controller platform.
Advanced Security	Enter the Brocade Mobility provided license key required to install the Role Based Firewall feature and increases the number of IPSec VPN tunnels. The number of IPSec tunnels varies by controller platform.
Advanced WIPS Licenses	Enter the Brocade Mobility provided license key required to install an advanced WIPS feature for client terminations and event sanctioning.

- Select **OK** to save the changes made to the applied licenses. Selecting **Reset** reverts the screen to its last saved configuration.

Assigning Certificates

A controller certificate links identity information with a public key enclosed in the certificate.

A *certificate authority* (CA) is a network authority that issues and manages security credentials and public keys for message encryption. The CA signs all digital certificates it issues with its own private key. The corresponding public key is contained within the certificate and is called a CA certificate. A browser must contain the CA certificate in its Trusted Root Library so it can trust certificates *signed* by the CA's private key.

Depending on the public key infrastructure, the digital certificate includes the owner's public key, the certificate expiration date, the owner's name and other public key owner information.

Each certificate is digitally signed by a *trustpoint*. The trustpoint signing the certificate can be a certificate authority, corporation or individual. A trustpoint represents a CA/identity pair containing the identity of the CA, CA-specific configuration parameters, and an association with an enrolled identity certificate.

SSH keys are a pair of cryptographic keys used to authenticate users instead of, or in addition to, a username/password. One key is private and the other is public key. *Secure Shell* (SSH) public key authentication can be used by a client to access managed resources, if properly configured. A RSA key pair must be generated on the client. The public portion of the key pair resides with the controller, while the private portion remains on a secure local area of the client.

To configure a controller's certificate usage:

1. Select the **Configuration** tab from the Web UI.
2. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

3. Select **Certificates** from the Device menu.

The screenshot displays the 'Management Security' configuration page. It is divided into three main sections: 'Management Security', 'RADIUS Security', and 'Information'.
 - Under 'Management Security', there are two items: 'HTTPS Trustpoint' and 'SSH RSA Key'. Each has a 'Pending' radio button selected and a text input field containing 'default-trustpoint' and 'default_rsa_key' respectively. Below each is a 'Stored' radio button, a dropdown menu (showing 'Device1' and 'Cert4'), and a 'Launch Manager' button.
 - Under 'RADIUS Security', there are two items: 'RADIUS Certificate Authority' and 'RADIUS Server Certificate'. Both have 'Pending' radio buttons selected and empty text input fields. Below each is a 'Stored' radio button, a dropdown menu (showing 'Device1'), and a 'Launch Manager' button.
 - The 'Information' section at the bottom contains a blue downward arrow icon and the text: '*Pending* Trustpoints and RSA Keys have not been verified to exist on the device.'
 - At the bottom right of the screen are three buttons: 'OK', 'Reset', and 'Exit'.

FIGURE 92 Device Certificates screen

- Set the following **Management Security** certificate configurations:

HTTPS Trustpoint	Either use the default-trustpoint or select the Stored radio button to enable a drop-down menu where an existing certificate/trustpoint can be leveraged. To leverage an existing device certificate for this device, select the Launch Manager button. For more information, see <i>Certificate Management on page 5-100</i> .
SSH RSA Key	Either use the default_rsa_key or select the Stored radio button to enable a drop-down menu where an existing certificate can be leveraged. To leverage an existing key for use with this target device, select the Launch Manager button. For more information, see <i>RSA Key Management on page 5-107</i> .

NOTE

Pending trustpoints and RSA keys are typically not verified as existing on a device.

- Set the following **RADIUS Security** certificate configurations:

RADIUS Certificate Authority	Either use the default-trustpoint or select the Stored radio button to enable a drop-down menu where an existing certificate can be leveraged. To leverage an existing certificate for this device, select the Launch Manager button.
RADIUS Server Certificate	Either use the default-trustpoint or select the Stored radio button to enable a drop-down menu where an existing certificate/trustpoint can be leveraged. To leverage an existing trustpoint for this device, select the Launch Manager button.

- Select **OK** to save the changes made to the certificate configurations. Selecting **Reset** reverts the screen to its last saved configuration.

For more information on the certification activities support by the controller, refer to the following:

- [Certificate Management](#)
- [RSA Key Management](#)
- [Certificate Creation](#)
- [Generating a Certificate Signing Request](#)

Certificate Management

Assigning Certificates

If not wanting to use an existing certificate or key with a selected device, an existing *stored* certificate can be leveraged from a different managed device for use with the target device. Device certificates can be imported and exported to and from the controller to a secure remote location for archive and retrieval as required for their application to other managed devices.

To configure trustpoints for use with certificates:

- Select **Launch Manager** from either the HTTPS Trustpoint, SSH RSA Key, RADIUS Certificate Authority or RADIUS Server Certificate parameters.

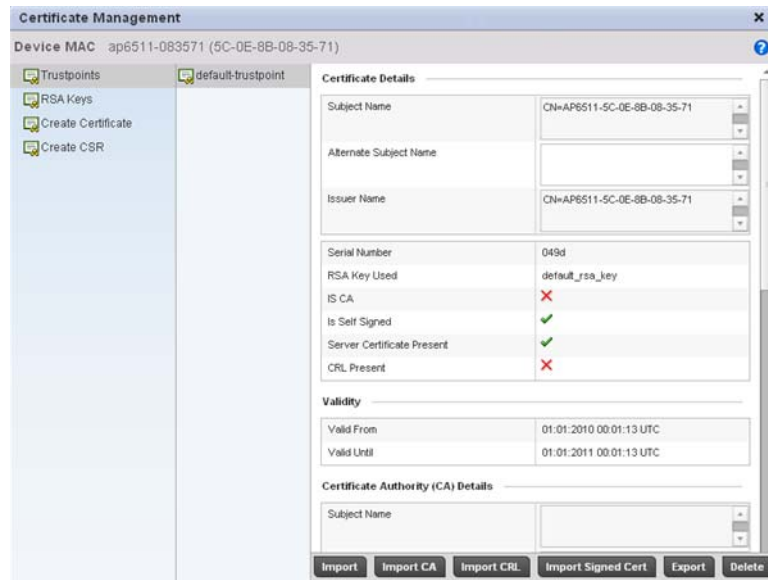


FIGURE 93 Certificate Management - Trustpoints screen

2. The Certificate Management screen displays with the **Trustpoints** portion displayed by default.
3. Select a device from amongst those displayed to review its certificate information.
4. Refer to the **Certificate Details** to review the certificate's properties, self-signed credentials, validity duration and CA information.
5. To optionally import a certificate to the controller, select the **Import** button from the Certificate Management screen.

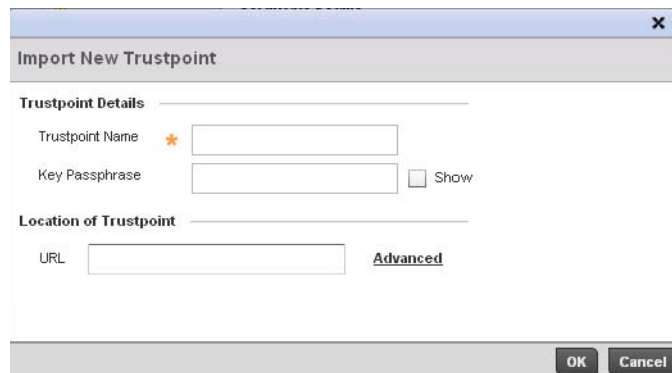


FIGURE 94 Certificate Management - Import New Trustpoint screen

6. Define the following configuration parameters required for the **Import** of the trustpoint.

Trustpoint Name	Enter the 32 character maximum name assigned to the target trustpoint. The trustpoint signing the certificate can be a certificate authority, corporation or individual.
Key Passphrase	Define the key used by both the controller and the server (or repository) of the target trustpoint. Select the Show textbox to expose the actual characters used in the key. Leaving the Show checkbox unselected displays the passphrase as a series of asterisks “*”.
URL	Provide the complete URL to the location of the trustpoint. If needed, select Advanced to expand the dialog to display network address information to the location of the target trustpoint. The number of additional fields that populate the screen is also dependent on the selected protocol.
Protocol	Select the protocol used for importing the target trustpoint. Available options include: tftp ftp sftp http cf usb1 usb2
Port	Use the spinner control to set the port. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> .
IP Address	Enter IP address of the server used to import the trustpoint This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> .
Hostname	Provide the hostname of the server used to import the trustpoint. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> .
Path	Specify the path to the trustpoint. Enter the complete relative path to the file on the server.

7. Select **OK** to import the defined trustpoint. Select **Cancel** to revert the screen to its last saved configuration.

8. To optionally import a CA certificate to the controller, select the **Import CA** button from the Certificate Management screen.

A CA is a network authority that issues and manages security credentials and public keys for message encryption. The CA signs all digital certificates it issues with its own private key. The corresponding public key is contained within the certificate and is called a CA certificate.

FIGURE 95 Certificate Management - Import CA Certificate screen

9. Define the following configuration parameters required for the **Import** of the CA certificate:

Trustpoint Name	Enter the 32 character maximum name assigned to the target trustpoint signing the certificate. A trustpoint represents a CA/identity pair containing the identity of the CA, CA-specific configuration parameters, and an association with an enrolled identity certificate.
Key Passphrase	Define the key used by both the controller and the server (or repository) of the target trustpoint. Select the Show textbox to expose the actual characters used in the key. Leaving the Show checkbox unselected displays the passphrase as a series of asterisks *-.*-.
URL	Provide the complete URL to the location of the trustpoint. If needed, select Advanced to expand the dialog to display network address information to the location of the target trustpoint. The number of additional fields that populate the screen is also dependent on the selected protocol.
Advanced / Basic	Click the Advanced or Basic link to switch between a basic URL and an advanced location to specify trustpoint location.
Protocol	Select the protocol used for importing the target CA certificate. Available options include: tftp ftp sftp http cf usb1 usb2
Port	Use the spinner control to set the port. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> .
IP Address	Enter IP address of the server used to import the CA certificate. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> .
Host	Provide the hostname of the server used to import the CA certificate. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> .
Path / File	Specify the path or filename of the CA certificate. Enter the complete relative path to the file on the server.

10. Select **OK** to import the defined CA certificate. Select **Cancel** to revert the screen to its last saved configuration.

11. To optionally import a CRL to the controller, select the **Import CRL** button from the Certificate Management screen.

If a certificate displays within the Certificate Management screen with a CRL, that CRL can be imported into the controller. A *certificate revocation list* (CRL) is a list of certificates that have been revoked or are no longer valid. A certificate can be revoked if the CA had improperly issued a certificate, or if a private-key is compromised. The most common reason for revocation is the user no longer being in sole possession of the private key.

12. For information on creating a CRL that can be used with a trustpoint, refer to [Setting the Certificate Revocation List \(CRL\) Configuration on page 7-370](#).

FIGURE 96 Certificate Management - Import CRL screen

Define the following configuration parameters required for the **Import** of the CRL

Trustpoint Name	Enter the 32 character maximum name assigned to the target trustpoint signing the certificate. A trustpoint represents a CA/identity pair containing the identity of the CA, CA-specific configuration parameters, and an association with an enrolled identity certificate.
From Network	Select the From Network radio button to provide network address information to the location of the target CRL. The number of additional fields that populate the screen is also dependent on the selected protocol. This is the default setting.
Cut and Paste	Select the Cut and Paste radio button to simply copy an existing CRL into the cut and past field. When pasting a CRL, no additional network address information is required.
URL	Provide the complete URL to the location of the CRL. If needed, select Advanced to expand the dialog to display network address information to the location of the CRL. The number of additional fields that populate the screen is also dependent on the selected protocol.
Protocol	Select the protocol used for importing the CRL. Available options include: tftp ftp sftp http cf usb1 usb2
Port	Use the spinner control to set the port. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
IP Address	Enter IP address of the server used to import the CRL. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Hostname	Provide the hostname of the server used to import the CRL. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Path	Specify the path to the CRL. Enter the complete relative path to the file on the server.

13. Select **OK** to import the CRL. Select **Cancel** to revert the screen to its last saved configuration.
14. To import a signed certificate to the controller, select the **Import Signed Cert** button from the Certificate Management screen.

Signed certificates (or root certificates) avoid the use of public or private CAs. A self-signed certificate is an identity certificate signed by its own creator, thus the certificate creator also signs off on its legitimacy. The lack of mistakes or corruption in the issuance of self signed certificates is central.

Self-signed certificates cannot be revoked which may allow an attacker who has already gained controller access to monitor and inject data into a connection to spoof an identity if a private key has been compromised. However, CAs have the ability to revoke a compromised certificate, preventing its further use.

FIGURE 97 Certificate Management - Import Signed Cert screen

15. Define the following configuration parameters required for the **Import** of the CA certificate:

Certificate Name	Enter the 32 character maximum trustpoint name with which the certificate should be associated.
From Network	Select the From Network radio button to provide network address information to the location of the signed certificate. The number of additional fields that populate the screen is also dependent on the selected protocol. This is the default setting.
Cut and Paste	Select the Cut and Paste radio button to simply copy an existing signed certificate into the cut and past field. When pasting a signed certificate, no additional network address information is required.
URL	Provide the complete URL to the location of the signed certificate. If needed, select Advanced to expand the dialog to display network address information to the location of the signed certificate. The number of additional fields that populate the screen is also dependent on the selected protocol.
Protocol	Select the protocol used for importing the target signed certificate. Available options include: tftp ftp sftp http cf usb1 usb2
Port	Use the spinner control to set the port. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> .

IP Address	Enter IP address of the server used to import the signed certificate. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Host	Provide the hostname of the server used to import the signed certificate. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Path / File	Specify the path to the signed certificate. Enter the complete relative path to the file on the server.

16. Select **OK** to import the signed certificate. Select **Cancel** to revert the screen to its last saved configuration

17. To optionally export a trustpoint from the controller to a remote location, select the **Export** button from the Certificate Management screen.

Once a certificate has been generated on the controller's authentication server, export the self signed certificate. A digital CA certificate is different from a self signed certificate. The CA certificate contains the public and private key pairs. The self certificate only contains a public key. Export the self certificate for publication on a Web server or file server for certificate deployment or export it in to an active directory group policy for automatic root certificate deployment.

18. Additionally export the key to a redundant RADIUS server so it can be imported without generating a second key. If there's more than one RADIUS authentication server, export the certificate and don't generate a second key unless you want to deploy two root certificates.

FIGURE 98 Certificate Management - Export Trustpoint screen

19. Define the following configuration parameters required for the **Export** of the trustpoint.

Trustpoint Name	Enter the 32 character maximum name assigned to the target trustpoint. The trustpoint signing the certificate can be a certificate authority, corporation or individual.
Key Passphrase	Define the key used by both the controller and the server (or repository) of the target trustpoint. Select the Show textbox to expose the actual characters used in the key. Leaving the Show checkbox unselected displays the passphrase as a series of asterisks "*" .
URL	Provide the complete URL to the location of the trustpoint. If needed, select Advanced to expand the dialog to display network address information to the location of the target trustpoint. The number of additional fields that populate the screen is also dependent on the selected protocol.

Protocol	Select the protocol used for exporting the target trustpoint. Available options include: tftp ftp sftp http cf usb1 usb2
Port	Use the spinner control to set the port. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> .
IP Address	Enter IP address of the server used to export the trustpoint. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Host	Provide the hostname of the server used to export the trustpoint. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Path / File	Specify the path to the trustpoint. Enter the complete relative path to the file on the server.

20. Select **OK** to export the defined trustpoint. Select **Cancel** to revert the screen to its last saved configuration.
21. To optionally delete a trustpoint, select the **Delete** button from within the Certificate Management screen. Provide the trustpoint name within the **Delete Trustpoint** screen and optionally select the **Delete RSA Key** checkbox to remove the RSA key along with the trustpoint. Select **OK** to proceed with the deletion, or **Cancel** to revert to the Certificate Management screen

RSA Key Management

[Assigning Certificates](#)

Refer to the RSA Keys screen to review existing RSA key configurations that have been applied to managed devices. If an existing key does not meet the needs of a pending certificate request, generate a new key or import/export an existing key to and from a remote location.

Rivest, Shamir, and Adleman (RSA) is an algorithm for public key cryptography. It's an algorithm that can be used for certificate signing and encryption. When a device trustpoint is created, the RSA key is the private key used with the trustpoint.

To review existing device RSA key configurations, generate additional keys or import/export keys to and from remote locations:

1. Select the **Launch Manager** button from either the SSH RSA Key, RADIUS Certificate Authority or RADIUS Server Certificate parameters (within the Certificate Management screen).
2. Select **RSA Keys** from the upper, left-hand, side of the Certificate Management screen.

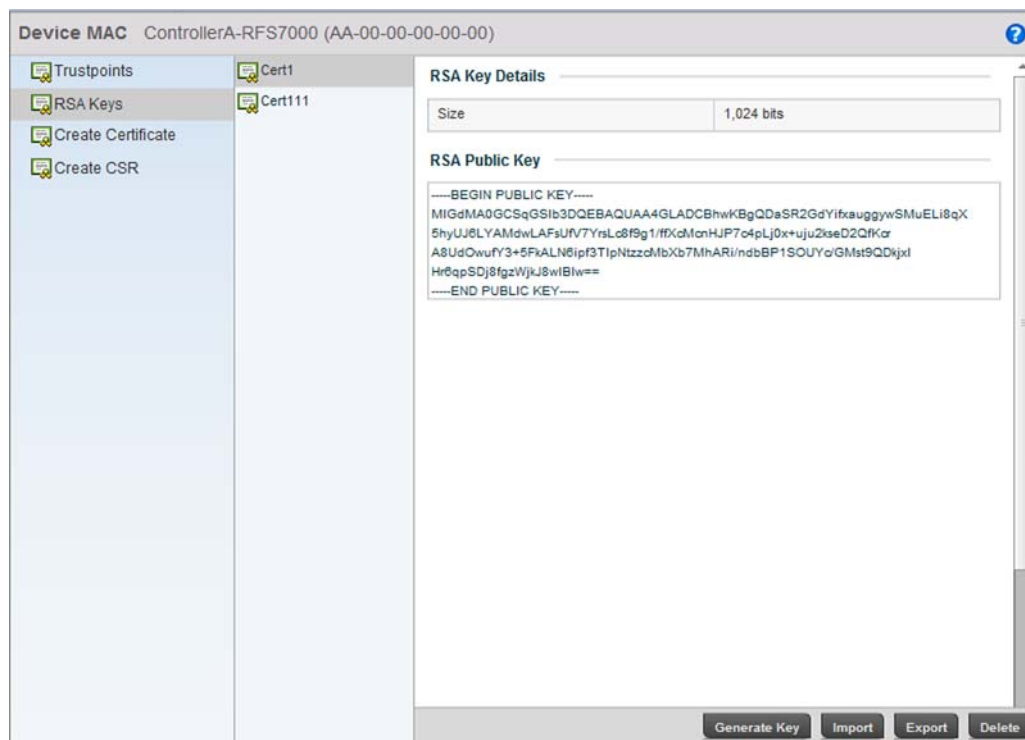


FIGURE 99 Certificate Management - RSA Keys screen

3. Select a listed device to review its current RSA key configuration.

Each key can have its size and character syntax displayed. Once reviewed, optionally generate a new RSA key, import a key from a selected device, export a key from the controller to a remote location or delete a key from a selected device.

4. Select **Generate Key** to create a new key with a defined size.



FIGURE 100 Certificate Management - Generate RSA Keys screen

5. Define the following configuration parameters required for the **Import** of the key:

Key Name	Enter the 32 character maximum name assigned to the RSA key.
Key Size	Use the spinner control to set the size of the key (between 1,024 - 2,048 bits). Brocade Mobility recommends leaving this value at the default setting of 1024 to ensure optimum functionality.

6. Select **OK** to generate the RSA key. Select **Cancel** to revert the screen to its last saved configuration.
7. To optionally import a CA certificate to the controller, select the **Import** button from the Certificate Management > RSA Keys screen.

FIGURE 101 Certificate Management - Import New RSA Key screen

8. Define the following configuration parameters required for the **Import** of the RSA key:

Key Name	Enter the 32 character maximum name assigned to identify the RSA key.
Key Passphrase	Define the key used by both the controller and the server (or repository) of the target RSA key. Select the Show textbox to expose the actual characters used in the passphrase. Leaving the Show checkbox unselected displays the passphrase as a series of asterisks “*”.
URL	Provide the complete URL to the location of the RSA key. If needed, select Advanced to expand the dialog to display network address information to the location of the target key. The number of additional fields that populate the screen is also dependent on the selected protocol.
Advanced / Basic	Click the Advanced or Basic link to switch between a basic URL and an advanced location to specify key location.
Protocol	Select the protocol used for importing the target key. Available options include: tftp ftp sftp http cf usb1 usb2

Port	Use the spinner control to set the port. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> .
IP Address	Enter IP address of the server used to import the RSA key. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> .
Host	Provide the hostname of the server used to import the RSA key. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> .
Path / File	Specify the path to the RSA key. Enter the complete relative path to the key on the server.

9. Select **OK** to import the defined RSA key. Select **Cancel** to revert the screen to its last saved configuration.
10. To optionally export a RSA key from the controller to a remote location, select the **Export** button from the Certificate Management > RSA Keys screen.

Export the key to a redundant RADIUS server so it can be imported without generating a second key. If there's more than one RADIUS authentication server, export the certificate and don't generate a second key unless you want to deploy two root certificates.

FIGURE 102 Certificate Management - Export RSA Key screen

11. Define the following configuration parameters required for the **Export** of the RSA key.

Key Name	Enter the 32 character maximum name assigned to the RSA key.
Key Passphrase	Define the key passphrase used by both the controller and the server. Select the Show textbox to expose the actual characters used in the passphrase. Leaving the Show checkbox unselected displays the passphrase as a series of asterisks “*”.
URL	Provide the complete URL to the location of the key. If needed, select Advanced to expand the dialog to display network address information to the location of the target key. The number of additional fields that populate the screen is also dependent on the selected protocol.

Protocol	Select the protocol used for exporting the RSA key. Available options include: tftp ftp sftp http cf usb1 usb2
Port	Use the spinner control to set the port. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> .
IP Address	Enter IP address of the server used to export the RSA key. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Host	Provide the hostname of the server used to export the RSA key. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Path / File	Specify the path to the key. Enter the complete relative path to the key on the server.

12. Select **OK** to export the defined RSA key. Select **Cancel** to revert the screen to its last saved configuration.
13. To optionally delete a key, select the **Delete** button from within the Certificate Management > RSA Keys screen. Provide the key name within the **Delete RSA Key** screen and select the **Delete Certificates** checkbox to remove the certificate the key supported. Select **OK** to proceed with the deletion, or **Cancel** to revert back to the Certificate Management screen.

Certificate Creation

Assigning Certificates

The Certificate Management screen provides the facility for creating new self-signed certificates. Self signed certificates (often referred to as root certificates) do not use public or private CAs. A self signed certificate is a certificate signed by its own creator, with the certificate creator responsible for its legitimacy.

To create a self-signed certificate that can be applied to a managed device:

1. Select the **Launch Manager** button from either the SSH RSA Key, RADIUS Certificate Authority or RADIUS Server Certificate parameters (within the Certificate Management screen).
2. Select **Create Certificate** from the upper, left-hand, side of the Certificate Management screen.

FIGURE 103 Certificate Management - Create Certificate screen

3. Define the following configuration parameters required to **Create New Self-Signed Certificate**:

- Certificate Name** Enter the 32 character maximum name assigned to identify the name of the trustpoint associated with the certificate. A trustpoint represents a CA/identity pair containing the identity of the CA, CA-specific configuration parameters, and an association with an enrolled identity certificate.
- Use an Existing RSA Key** Select the radio button and use the drop-down menu to select the existing key used by both the controller and the server (or repository) of the target RSA key.
- Create a New RSA Key** To create a new RSA key, select the radio button to define a 32 character name used to identify the RSA key. Use the spinner control to set the size of the key (between 1,024 - 2,048 bits). Brocade Mobility recommends leaving this value at the default setting of 1024 to ensure optimum functionality. For more information on creating a new RSA key, see *RSA Key Management on page 5-107*.

4. Set the following **Certificate Subject Name** parameters required for the creation of the certificate:

- Certificate Subject Name** Select either the *auto-generate* radio button to automatically create the certificate's subject credentials or select *user-defined* to manually enter the credentials of the self signed certificate. The default setting is auto-generate.
- Country (C)** Define the Country used in the certificate. The field can be modified by the user to other values. This is a required field and must not exceed 2 characters.
- State (ST)** Enter a State/Prov. for the state or province name used in the certificate. This is a required field.
- City (L)** Enter a City to represent the city name used in the certificate. This is a required field.

- Organization (O)** Define an Organization for the organization used in the certificate. This is a required field.
- Organizational Unit (OU)** Enter an Org. Unit for the name of the organization unit used in the certificate. This is a required field.
- Common Name (CN)** If there's a common name (IP address) for the organizational unit issuing the certificate, enter it here.

5. Select the following **Additional Credentials** required for the generation of the self signed certificate:

- Email Address** Provide an email address used as the contact address for issues relating to this certificate request.
- Domain Name)** Enter a *fully qualified domain name* (FQDN) is an unambiguous domain name that specifies the node's position in the DNS tree hierarchy absolutely. To distinguish an FQDN from a regular domain name, a trailing period is added. For example, *somehost.example.com*. An FQDN differs from a regular domain name by its absoluteness, since a suffix is not added.
- IP Address** Specify the IP address used as the controller destination for certificate requests.

6. Select the **Generate Certificate** button at the bottom of the Certificate Management > Create Certificate screen to produce the certificate.

Generating a Certificate Signing Request

[Assigning Certificates](#)

A *certificate signing request* (CSR) is a request to a certificate authority to apply for a digital identity certificate. The CSR is a block of encrypted text generated on the server the certificate is used on. It contains the organization name, common name (domain name), locality and country.

A RSA key must be either created or applied to the certificate request before the certificate can be generated. A private key is not included in the CSR, but is used to digitally sign the completed request. The certificate created with a particular CSR only worked with the private key generated with it. If the private key is lost, the certificate is no longer functional. The CSR can be accompanied by other identity credentials required by the certificate authority, and the certificate authority maintains the right to contact the applicant for additional information.

If the request is successful, the CA sends an identity certificate digitally signed with the private key of the CA.

To create a CSR:

1. Select the **Launch Manager** button from either the SSH RSA Key, RADIUS Certificate Authority or RADIUS Server Certificate parameters (within the Certificate Management screen).
2. Select **Create CSR** from the upper, left-hand, side of the Certificate Management screen.

FIGURE 104 Certificate Management - Create CSR screen

3. Define the following configuration parameters required to **Create New Certificate Signing Request (CSR)**:

Use an Existing RSA Key Select the radio button and use the drop-down menu to set the key used by both the controller and the server (or repository) of the target RSA key.

Create a New RSA Key To create a new RSA key, select the radio button to define 32 character name used to identify the RSA key. Use the spinner control to set the size of the key (between 1,024 - 2,048 bits). Brocade Mobility recommends leaving this value at the default setting of 1024 to ensure optimum functionality. For more information on creating a new RSA key, see *RSA Key Management on page 5-107*.

4. Set the following **Certificate Subject Name** parameters required for the creation of the certificate:

Certificate Subject Name Select either the *auto-generate* radio button to automatically create the certificate's subject credentials or select *user-defined* to manually enter the credentials of the self signed certificate. The default setting is auto-generate.

Country (C) Define the Country used in the CSR. The field can be modified by the user to other values. This is a required field and must not exceed 2 characters.

State (ST) Enter a State/Prov. for the state or province name used in the CSR. This is a required field.

City (L) Enter a City to represent the city name used in the CSR. This is a required field.

Organization (O)	Define an Organization for the organization used in the CSR. This is a required field.
Organizational Unit (OU)	Enter an Org. Unit for the name of the organization unit used in the CSR. This is a required field.
Common Name (CN)	If there's a common name (IP address) for the organizational unit issuing the certificate, enter it here.

5. Select the following **Additional Credentials** required for the generation of the CSR:

Email Address	Provide an email address used as the contact address for issues relating to this CSR.
Domain Name)	Enter a <i>fully qualified domain name</i> (FQDN) is an unambiguous domain name that specifies the node's position in the DNS tree hierarchy absolutely. To distinguish an FQDN from a regular domain name, a trailing period is added. ex: somehost.example.com. An FQDN differs from a regular domain name by its absoluteness; as a suffix is not added.
IP Address	Specify the controller IP address used as the controller destination for certificate requests.

6. Select the **Generate CSR** button to produce the CSR.

RF Domain Overrides

Use **RF Domain Overrides** to define configurations overriding the configuration set by the target device's original RF Domain assignment.

RF Domains allow administrators to assign configuration data to multiple devices deployed in a common coverage area (floor, building or site). In such instances, there's many configuration attributes these devices share, since their general client support roles are quite similar. However, device configurations may need periodic refinement from their original RF Domain administered design.

A controller configuration contains (at a minimum) one default RF Domain, but can optionally use additional user defined RF Domains:

- *Default RF Domain* - Automatically assigned to each controllers and associated Access Points by default. A default RF Domain is unique to a specific controller RFS4000, RFS6000 or RFS7000 or access point (Brocade Mobility 650 Access Point, Brocade Mobility 6511 Access Point, or Brocade Mobility 7131 Access Point) model.
- *User Defined RF Domains* - Created by administrators and manually assigned to individual controllers or Access Points, but can be automatically assigned to Access Points using adoption policies.

Each controller and Access Point is assigned only one RF Domain at a time. However, a user defined RF Domain can be assigned to multiple controllers or Access Points as required. User defined RF Domains can be manually assigned to controllers and Access Points or automatically assigned to Access Points using an auto provisioning policy. The more devices assigned a single RF Domain, the greater the likelihood that one of the device's configurations will require an override that deviates that device's configuration from the original RF Domain assignment shared by the others.

To review the RF Domain's original configuration requirements and the options available for a target device, refer to [Managing RF Domains on page 8-408](#).

To define a device's RF Domain override configuration:

1. Select the **Configuration** tab from the Web UI.
2. Select **Devices** from the Configuration tab.
3. The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.
4. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

5. Expand the **RF Domain Overrides** menu option to display its sub-menu options.
6. Select **RF Domain**.

FIGURE 105 RF Domain Overrides screen

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

7. Refer to the **Basic Configuration** field to review the basic settings defined for the target device's RF Domain configuration, and optionally assign/remove overrides to and from specific parameters.

Location	Displays the location set for the device as part of its RF Domain configuration.
Contact	Displays the contact information set for the device as part of its RF Domain configuration.
Time Zone	Displays the time zone set for the device as part of its RF Domain configuration.
Country Code	Displays the country code set for the device as part of its RF Domain configuration.
VLAN for Control Traffic	Displays the VLAN for Control Traffic setting for the device as part of its RF Domain configuration.

8. Refer to the **Smart RF** section to configure Smart RF policy and dynamic channel settings.

Smart RF Policy	Use the Smart RF Policy drop-down menu to apply a Smart RF policy to the RF Domain. When a radio fails or is faulty, a Smart RF policy can provide automatic recovery by instructing neighboring Access Point radios to increase their transmit power to compensate for the coverage loss.
Enable Dynamic Channel	Check this box to enable dynamic channel switching for Smart RF radios.
2.4 GHz Channels	Select channels from the pull-down menu and click the down arrow to move it to the list of channels used for 2.4GHz Smart RF radios.
5 GHz Channels	Select channels from the pull-down menu and click the down arrow to move it to the list of channels used for 5GHz Smart RF radios.
2.4 GHz Radios	Select radios from the drop-down menu and click the down arrow to move it to the list of channels used for 2.4GHz Smart RF radios.
5 GHz Radios	Select radios from the drop-down menu and click the down arrow to move it to the list of channels used for 5GHz Smart RF radios.

9. Select the **Create** icon to define a new Smart RF policy that can be applied to the RF Domain, or select the **Edit** icon to modify or override an existing Smart RF policy.

For an overview of Smart RF and instructions on how to create a Smart RF policy that can be used with a RF Domain, see [Smart RF Policy on page 6-301](#).

10. Use the **WIPS Policy** drop-down menu to apply a WIPS policy to the RF Domain.

The wireless controller supports the *Wireless Intrusion Protection System* (WIPS) to provide continuous protection against wireless threats and act as an additional layer of security complementing wireless VPNs and encryption and authentication policies. The wireless controller supports WIPS through the use of dedicated sensor devices designed to actively detect and locate unauthorized AP devices. After detection, they use mitigation techniques to block devices using manual termination, air lockdown or port suppression.

11. Select the **Create** icon to define a new WIPS policy that can be applied to the RF Domain, or select the **Edit** icon to modify or override an existing WIPS policy.

For an overview of WIPS and instructions on how to create a WIPS policy that can be used with a RF Domain, see [Intrusion Prevention on page 9-446](#).

12. Refer to the **Statistics** field to set the following data:

- NoC Update Interval** Set a NoC Update interval of 0, or between 5-300 seconds for updates from the RF Domain manager to the controller.
- Window Index** Use the spinner control to set a numerical index used as an identifier for each RF Domain statistics defined.
- Sample Interval** Use the spinner control to define the interval (in seconds) used by the controller to capture windowed statistics supporting the listed RF Domain configuration. The default is 5 seconds.
- Window Size** Use the spinner control to set the number of samples used by the controller to define RF Domain statistics. The default value is 6 samples.

13. Select **OK** to save the changes and overrides made to the RF Domain configuration. Selecting **Reset** reverts the screen to its last saved configuration.

14. Select **Sensor Configuration** from within the expanded RF Domain Overrides menu.

Server Id	IP Address	Port	
1	157.96.98.42	443	

Add Row

Note: The default port used by AirDefense Server is 443 and the default port used by Advanced-WIPS on a Controller is 8443

OK **Reset** **Exit**

FIGURE 106 Sensor Appliance Configuration Override screen

15. Define a **Sensor Appliance Configuration** for dedicating a WIPS server resource for client terminations and WIPS event logging.

16. Optionally set up to 3 overrides for the listed device's sensor server assignment:

Server Id	Use the spinner control set a numerical index for the sensor server to differentiate it from other servers. Up to 3 sensor server resources can be defined. Select the Add Row + button as needed to add additional servers.
IP Address	Set the IP addresses of up to 3 sensor servers for supporting WIPS events on behalf of the selected device.
Port	Assign the port number of the sensor server using the spinner control. The default port is port 443.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

17. Select **OK** to save the changes and overrides made to the Sensor Appliance Configuration. Selecting **Reset** reverts the screen to its last saved configuration.

18. Select **WLAN Override** from within the expanded RF Domain Overrides.

The screenshot shows the 'WLAN Override' screen with the 'Override SSID' tab selected. The table below is a representation of the data shown in the screenshot:

WLAN	SSID	
vLAN101	eng2	

Buttons: Add Row, OK, Reset, Exit

FIGURE 107 WLAN Override screen - Override SSID tab

The WLAN Override screen displays with the **Override SSID** tab displayed by default.

19. Optionally define up to 3 overrides for the listed device's WLAN SSID assignment:

WLAN Optionally use the drop-down menu to change the WLAN assignment for the listed device. Select either the **Create** icon to define a new WLAN's configuration, or select the **Edit** icon to modify an existing WLAN configuration. For additional information on either creating or editing a WLAN's configuration, see *Basic WLAN Configuration on page 6-229*.

SSID Optionally change the SSID associated with the WLAN. The WLAN name is auto-generated using the SSID until changed (overridden). The maximum number of characters used for the SSID is 32.

20. Select the **Add Row +** button as needed to add additional WLAN SSID overrides.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the Basic Configuration section of the device and click the Clear Overrides button. This will remove all overrides from the device.

21. Select **OK** to save the changes and overrides. Selecting **Reset** reverts the screen to its last saved configuration.

22. Select the **Override VLAN** tab to review any VLAN assignment overrides that may have been or optionally add or edit override configurations.




FIGURE 108 WLAN Override screen - Override VLAN tab

The Override VLANs tab displays the VLANs assigned to the WLAN on the device. Select **Add** to create a new client limit configuration for a specific WLAN and VLAN or **Edit** to modify an existing configuration.

23. Optionally define a VLAN's wireless client limit override configuration.

WLAN	If adding a new VLAN client limit assignment, select the target WLAN from the drop-down menu. Select the Create icon to create a new controller WLAN configuration, or select the Edit icon to modify an existing controller WLAN. Refer to <i>Wireless LAN Policy on page 6-228</i> to define a new controller WLAN that can be applied to a managed device.
VLANS	Use the spinner control to set a VLAN ID (between 1 - 4094). Refer to Basic WLAN Configuration on page 6-229 to define whether a single VLAN or VLAN pool is used with a WLAN.
Wireless Client Limit	Use the spinner control to set the client limit on the listed VLAN. The default setting is 0, implying the maximum client count is unlimited.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

24. Select **OK** to save the changes and overrides. Selecting **Reset** reverts the screen to its last saved configuration.

Profile Overrides

Profiles enable administrators to assign a common set of configuration parameters and policies to controllers and Access Points. Profiles can be used to assign shared or *unique* network, wireless and security parameters to wireless controllers and Access Points across a large, multi segment, site. The configuration parameters within a profile are based on the hardware model the profile was created to support. The controller supports both default and user defined profiles implementing new features or updating existing parameters to groups of Wireless Controllers or Access Points. The central benefit of a profile is its ability to update devices collectively without having to modify individual device configurations. Power and Adoption overrides apply specifically to Access Points, while Cluster configuration overrides apply to only controller configurations.

However, device profile configurations may need periodic refinement from their original administered design. Consequently, a device profile could require modification from a profile configuration shared amongst numerous devices deployed within a particular site.

Use Profile Overrides to define configurations overriding the parameters set by the target device's original profile assignment.

To review a profile's original configuration requirements and the options available for a target device, refer to [General Profile Configuration on page 7-316](#).

To define a device's general profile override configuration:

1. Select the **Configuration** tab from the Web UI.
2. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

3. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

4. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
5. Select **General** if it doesn't display by default.

The screenshot displays the 'Settings' window for 'Profile Overrides - General'. It is organized into several sections:

- Settings:** Includes 'IP Routing' with an information icon and a checked checkbox.
- Auto-Provisioning Policy:** Features a dropdown menu for 'Auto-Provisioning Policy', an information icon, and a checked checkbox for 'Learn and save network configuration'.
- AP300 Auto-Provisioning Policy:** Contains 'Adopt Unknown APs Automatically' (checked), 'Allowed List' (with a dropdown set to '<none>' and a list containing 'AA-00-22-33-44-55'), and 'Deny List' (with a dropdown set to '<none>' and a list containing 'AA-23-55-89-88-11').
- Network Time Protocol (NTP):** A table with columns: Server IP, Authentication Key, Prefer, Autokey, Key, Version, and a delete icon. The table is currently empty.

At the bottom right of the table is an 'Add Row' button. At the very bottom of the window are 'OK', 'Reset', and 'Exit' buttons.

FIGURE 109 Profile Overrides - General screen

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

6. In the **Settings** section check the **IP Routing** checkbox to enable routing for the device.
7. Refer to the **Auto Provisioning Policy** section to select an **Auto Provisioning Policy** or create a new one.

Auto Provisioning Policy Select an Auto Provisioning Policy from the pulldown menu. To create a new Auto Provisioning Policy click the create icon. For more information on creating an auto provisioning policy that can be applied to a controller profile, see *Auto Provisioning Policies* on page 5-216.

Learn and save network configuration Check the Learn and save network configuration checkbox to enable the device to learn and save network information.

8. Select **+ Add Row** below the **Network Time Protocol (NTP)** table to define (or override) the configurations of NTP server resources the controller uses it obtain its system time. Set the following parameters to define the NTP configuration:

Server IP	Set the IP address of each server added as a potential NTP resource.
Authentication Key	Select the number of the associated Authentication Key for the NTP resource.
Prefer	Select the radio button to designate this particular NTP resource as preferred. If using multiple NTP resources, preferred resources will be given first opportunity to connect to the controller and provide NTP calibration.
AutoKey	Select the radio button to enable an autokey configuration for the controller and NTP resource. The default setting is disabled.
Key	If an autokey is not being used, manually enter a 64 character maximum key the controller and NTP resource share to securely interoperate.
Version	Use the spinner control to specify the version number used by this NTP server resource. The default setting is 0.

9. Select **OK** to save the changes and overrides made to the general profile configuration. Select **Reset** to revert to the last saved configuration.

Controller Cluster Configuration Overrides (Controllers Only)

A redundancy group (cluster) is a set of controllers (nodes) uniquely defined by a controllers profile configuration. Within the redundancy group, members discover and establish connections to other controller members and provide wireless network self-healing support in the event of cluster member failure.

A cluster's AP load balance is typically distributed evenly amongst the controllers in the cluster. Define how often this profile is load balanced for AP radio distribution as often as you feel required, as radios can come and go and controller members can join and exit the cluster. For information on setting a profile's original cluster configuration (before applying an override), see [Profile Cluster Configuration \(Controllers Only\) on page 7-318](#).

As cluster memberships increase or decrease and their load requirements change, a controller's profile may need an override applied to best suit a site's cluster requirements.

NOTE

There is a limit of 2 controllers that can be configured in a cluster.

To apply an override (if required) to a controller profile cluster configuration:

1. Select the **Configuration** tab from the Web UI.
2. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.
3. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.
4. Select **Profile Overrides** from the Device menu to expand it into sub menu options.

5. Select **Cluster**.**NOTE**

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

The screenshot displays the 'Cluster Settings' and 'Cluster Member' configuration sections. The 'Cluster Settings' section includes:

- Cluster Mode:** Radio buttons for 'Active' (selected) and 'Standby'.
- Cluster Name:** Text input field containing 'cm4'.
- Master Priority:** Spin-down menu set to '128' (range: 1 to 255).
- Handle STP Convergence:** Checkmark icon.
- Force Configured State:** Checkmark icon.
- Force Configured State Delay:** Spin-down menu set to '5' (range: 3 to 1,800 minutes).

 The 'Cluster Member' section includes:

- Cluster VLAN:** Spin-down menu set to '1' (range: 1 to 4,094).
- Table:** A table with columns 'Member IP Address', 'Routing Level', and a delete icon. It contains one empty row.
- Add Row:** A button at the bottom right of the table.

 At the bottom of the screen, there are buttons for 'Restore Configured State', 'Force Active', 'Force Standby', 'OK', 'Reset', and 'Exit'.

FIGURE 110 Profile Overrides - Controller Cluster screen

6. Optionally define the following **Cluster Settings** and overrides:**Cluster Mode**

A member can be in either an *Active* or *Standby* mode. All active member controllers can adopt Access Points. Standby members only adopt Access Points when an active member has failed or sees an Access Point not adopted by a controller. The default cluster mode is Active and enabled for use with the controller profile.

Cluster Name

Define a name for the cluster name unique to its configuration or profile support requirements. The name cannot exceed 64 characters.

Master Priority

Set a priority value between 1 and 255 with the higher value being given higher priority. This configuration is the device's priority to become cluster master. In cluster environment one device from cluster members is elected as cluster master. This configuration is the device's priority to become cluster master. The default value is 128.

Handle STP Convergence	Select the radio button to enable <i>Spanning Tree Protocol</i> (STP) convergence for the controller. In general, this protocol is enabled in layer 2 networks to prevent network looping. Spanning Tree is a network layer protocol that ensures a loop-free topology in a mesh network of inter-connected layer 2 controllers. The spanning tree protocol disables redundant connections and uses the least costly path to maintain a connection between any two controllers in the network. If enabled, the network forwards data only after STP convergence. Enabling STP convergence delays the redundancy state machine execution until the STP convergence is completed (the standard protocol value for STP convergence is 50 seconds). Delaying the state machine is important to load balance APs at startup. The default setting is disabled.
Force Configured State	Select the radio button to allow this controller to take over for an active controller member if it were to fail. A standby controller in the cluster takes over APs adopted by the failed active controller. If the failed active controller were to come back up, the active controller starts a timer based on the Auto Revert Delay interval. At the expiration of the Auto Revert Delay, the standby controller releases all adopted APs and goes back to a monitoring mode. The Auto Revert Delay timer is stopped and restarted if the active controller goes down and comes up during the Auto Revert Delay interval. The default value is disabled.
Force Configured State Delay	Specify a delay interval in minutes (1 - 1,800). This is the interval a standby controller waits before releasing adopted APs and goes back to a monitoring mode when an active controller becomes active again after a failure. The default interval is 5 seconds.

7. Within the **Cluster Member** field, select the **Cluster VLAN** checkbox to enable a spinner control to designate the controller VLAN where cluster members are reachable. Specify a VLAN in the range of 1 - 4094.
Specify the IP addresses of the VLAN's cluster members using the IP Address table.
8. Select an **Auto-Provisioning Policy** from the pulldown menu. To create a new Auto-Provisioning Policy click the create icon.
9. Select **OK** to save the changes and overrides made to the profile's cluster configuration. Select **Reset** to revert to the last saved configuration.

Access Point Adoption Overrides (Access Points Only)

Adoption is the process an Access Point uses to discover controllers available in the network, pick the most desirable controller, establish an association with the controller and optionally obtain an image upgrade, obtains its configuration and considers itself provisioned. This is a configurable activity that can be supported within a device profile and applied to other Access Points supported by the profile. Individual attributes of an Access Point profile auto provisioning policy can be overridden as specific parameters require modification.

At adoption, an Access Point solicits and receives multiple adoption responses from controllers available on the network. These adoption responses contain loading policy information the Access Point uses to select the optimum controller for adoption. By default, an auto provisioning policy generally distributes AP adoption evenly amongst available controllers. Modify existing adoption policies or create a new one as needed to meet the adoption requirements of a device and their assigned controller profile.

NOTE

A device configuration does not need to be present for an auto provisioning policy to take effect. Once adopted, and the device's configuration is defined and applied by the controller, the auto provisioning policy mapping does not have impact on subsequent adoptions by the same device.

An auto provisioning policy enables an administrator to define adoption rules for the supported Brocade Mobility Access Points capable of being adopted by a wireless controller.

To define an Access Point's adoption configuration or apply an override:

1. Select the Devices from the Web UI.
2. Select **Profiles** from the Configuration tab.
3. Select **Profile Overrides** to expand its sub-menu items.
4. Select **Adoption**.

A screen displays where an Access Point's adoption configuration can be defined and overridden for a controller profile.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

Controller Group

Preferred Group

Controller VLAN

VLAN (1 to 4,094)

Controller Hostnames

Host	Pool	Routing Level	

Add Row

OK Reset Exit

FIGURE 111 Access Point Adoption Override screen

5. Define or override the **Preferred Group** used as optimal group of controllers for the Access Point's adoption. The name of the preferred group cannot exceed 64 characters.

6. Select the checkbox to define or override a **VLAN** the Access Point's associating controller is reachable on.

VLANs 0 and 4,095 are reserved and cannot be used by a controller VLAN.

7. Enter **Controller Hostnames** as needed to define or override controller resources for Access Point adoption.

Select **+ Add Row** as needed to populate the table with IP Addresses or Hostnames of controllers used as Access Point adoption resources into the managed network.

Host	Use the drop-down menu to specify whether the controller adoption resource is defined as a (non DNS) IP Address or a Hostname . Once defined, provide the numerical IP or Hostname. A Hostname cannot exceed 64 characters.
Pool	Use the spinner controller to set a pool of either 1 or 2. This is the pool the target controller belongs to.
Remote	Select the checkbox if the controller IP address or hostname provided within the host field resides within a remote RF Domain. This setting is enabled by default.

8. Select **OK** to save the changes and overrides made to the Access Point profile adoption configuration. Select **Reset** to revert to the last saved configuration.

Access Point Radio Power Overrides (Access Points Only)

A controller profile can manage the transmit output power of the Access Point radios it supports within the managed network.

NOTE

The Power option only appears within the Profile Overrides menu tree if an Access Point is selected from within the main Devices screen. Power management is configured differently for controllers, so the Power screen does not display for RFS6000 and RFS7000 model controllers or the Brocade Mobility 650 Access Point. It only displays on Brocade Mobility 6511 Access Point, Brocade Mobility 7131 Access Point models.

Use the Power screen to set or override one of two power modes (3af or Auto) for a managed Access Point. When automatic is selected, the Access Point safely operates within available power. Once the power configuration is determined, the Access Point configures its operating power characteristics based on its model and power configuration.

An Access Point uses a *complex programmable logic device* (CPLD). The CPLD determines proper supply sequencing, the maximum power available and other status information. One of the primary functions of the CPLD is to determine the Access Point's maximum power budget. When an Access Point is powered on (or performing a cold reset), the CPLD determines the maximum power provided by the POE device and the budget available to the Access Point. The CPLD also determines the access point hardware SKU and the number of radios. If the Access Point's POE resource cannot provide sufficient power to run the access point (with all intended interfaces enabled), some of the following interfaces could be disabled or modified:

- The Access Point's transmit and receive algorithms could be negatively impacted
- The Access Point's transmit power could be reduced due to insufficient power
- The Access Point's WAN port configuration could be changed (either enabled or disabled)

To define an Access Point's power configuration or apply an override to an existing parameter:

1. Select the Devices tab from the Web UI.

2. Select **Profile Overrides** to expand its sub menu items.
3. Select **Power**.

A screen displays where an Access Point's power configuration can be defined or overridden for a controller profile.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

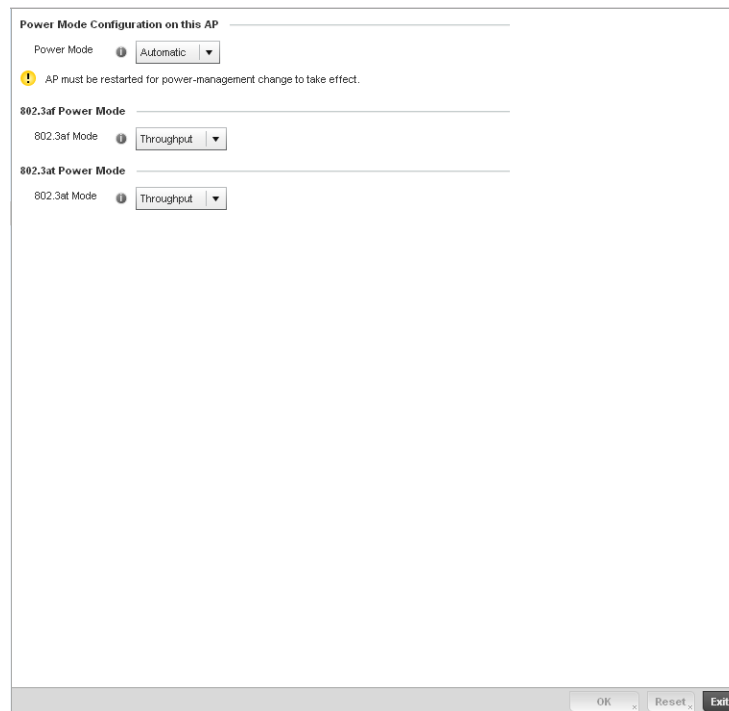


FIGURE 112 Access Point Profile Power Override screen

4. Use the **Power Mode** drop-down menu to set or override the **Power Mode Configuration on this AP**.

NOTE

Single radio model Access Point's always operate using a full power configuration. The power management configurations described in this section do not apply to single radio models.

When an Access Point is powered on for the first time, the system determines the power budget available to the Access Point. Using the **Automatic** setting, the Access Point automatically determines the best power configuration based on the available power budget. Automatic is the default setting.

If 802.3af is selected, the Access Point assumes 12.95 watts are available. If the mode is changed, the Access Point requires a reset to implement the change. If 802.3at is selected, the Access Point assumes 23 - 26 watts are available.

5. Set or override the Access Point radio's **802.3af Power Mode** and the radio's **802.3at Power Mode**.

Use the drop-down menu to define a mode of either **Range** or **Throughput**.

Select **Throughput** to transmit packets at the radio's highest defined basic rate (based on the radio's current basic rate settings). This option is optimal in environments where the transmission range is secondary to broadcast/multicast transmission performance. Select **Range** when range is preferred over performance for broadcast/multicast (group) traffic. The data rates used for range are the lowest defined basic rates. Throughput is the default setting for both 802.3af and 802.3at.

6. Select **OK** to save the changes and overrides made to the Access Point power configuration. Select **Reset** to revert to the last saved configuration.

Profile Interface Override Configuration

A controller profile's interface configuration can be defined to support separate physical Ethernet configurations both unique and specific to RFS4000, RFS4000 and RFS7000 model controllers. Ports vary depending on controller platform, but controller models do have some of the same physical interfaces.

A controller requires its Virtual Interface be configured for layer 3 (IP) access or layer 3 service on a VLAN. A controller's Virtual Interface defines which IP address is associated with each VLAN ID the controller is connected to.

If the profile is configured to support an Access Point radio, an additional Radios option is available, unique to the Access Point's radio configuration.

Each profile interface configuration can have overrides applied to customize the configuration to a unique controller deployment. However, once an override is applied to this configuration it becomes independent from the profile that may be shared by a group of devices in a specific deployment and may need careful administration until a profile can be re-applied to the target controller. For more information, refer to the following:

- [Ethernet Port Override Configuration](#)
- [Virtual Interface Override Configuration](#)
- [Port Channel Override Configuration](#)
- [Radio Override Configuration](#)

Ethernet Port Override Configuration

Profile Interface Override Configuration

The ports available on a controller vary depending on the platform. The following ports are available on RFS4000, RFS6000 and RFS7000 model controllers:

- RFS4000 - ge1, ge2, ge3, ge4, ge5, up1
- RFS6000 - ge1, ge2, ge3, ge4, ge5, ge6, ge7, ge8, me1, up1
- RFS7000 - ge1, ge2, ge3, ge4, me1

GE ports are available on the RFS4000, RFS6000 and RFS7000 platforms. GE ports on the RFS4000 and RFS6000 are RJ-45 supporting 10/100/1000Mbps. GE ports on the RFS7000 can be RJ-45 or fiber ports supporting 10/100/1000Mbps.

ME ports are available on RFS6000 and RFS7000 platforms. ME ports are out-of-band management ports used to manage the controller via CLI or Web UI, even when the other ports on the controller are unreachable.

7. Refer to the following to assess port status and performance:

Name	Displays the physical controller port name reporting runtime data and statistics. Supported ports vary depending on controller model. RFS4000 - ge1, ge2, ge3, ge4, ge5, up1 RFS6000 - ge1, ge2, ge3, ge4, ge5, ge6, ge7, ge8, me1, up1 RFS7000 - ge1, ge2, ge3, ge4, me1
Type	Displays the physical controller port type. <i>Cooper</i> is used on RJ45 Ethernet ports and <i>Optical</i> materials are used on fiber optic gigabit Ethernet ports.
Description	Displays an administrator defined description for each listed controller port.
Admin Status	A green checkmark defines the port as active and currently enabled with the controller profile. A red "X" defines the port as currently disabled and not available for use. The interface status can be modified with the port configuration as needed.
Mode	Displays the profile's switching mode as either <i>Access</i> or <i>Trunk</i> (as defined within the Ethernet Port Basic Configuration screen). If <i>Access</i> is selected, the listed port accepts packets only from the native VLAN. Frames are forwarded untagged with no 802.1Q header. All frames received on the port are expected as untagged and mapped to the native VLAN. If set to <i>Trunk</i> , the port allows packets from a list of VLANs added to the trunk. A port configured as <i>Trunk</i> supports multiple 802.1Q tagged VLANs and one Native VLAN which can be tagged or untagged.
Native VLAN	Lists the numerical VLAN ID (1 - 4094) set for the native VLAN. The native VLAN allows an Ethernet device to associate untagged frames to a VLAN when no 802.1Q frame is included in the frame. Additionally, the native VLAN is the VLAN untagged traffic is directed over when using a port in trunk mode.
Tag Native VLAN	A green checkmark defines the native VLAN as tagged. A red "X" defines the native VLAN as untagged. When a frame is tagged, the 12 bit frame VLAN ID is added to the 802.1Q header so upstream Ethernet devices know which VLAN ID the frame belongs to. The device reads the 12 bit VLAN ID and forwards the frame to the appropriate VLAN. When a frame is received with no 802.1Q header, the upstream device classifies the frame using the default or native VLAN assigned to the Trunk port. A native VLAN allows an Ethernet device to associate untagged frames to a VLAN when no 802.1Q frame is included in the frame.
Allowed VLANs	Displays those VLANs allowed to send packets over the listed controller port. Allowed VLANs are only listed when the mode has been set to <i>Trunk</i> .

8. To edit or override the configuration of an existing controller port, select it from amongst those displayed and select the **Edit** button. The Ethernet port **Basic Configuration** screen displays by default.

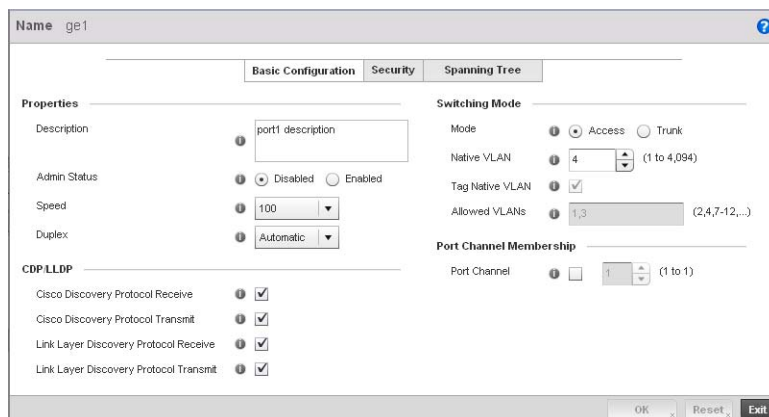


FIGURE 114 Profile Overrides - Ethernet Ports Basic Configuration screen

9. Set or override the following Ethernet port **Properties**:

- | | |
|---------------------|--|
| Description | Enter a brief description for the controller port (64 characters maximum). The description should reflect the port's intended function to differentiate it from others with similar configurations or perhaps just the name of the physical port. |
| Admin Status | Select the Enabled radio button to define this port as active to the controller profile it supports. Select the Disabled radio button to disable this physical controller port in the controller profile. It can be activated at any future time when needed. |
| Speed | Select the speed at which the port can receive and transmit the data. Select either 10 Mbps, 100 Mbps, 1000 Mbps. Select either of these options to establish a 10, 100 or 1000 Mbps data transfer rate for the selected half duplex or full duplex transmission over the port. These options are not available if Auto is selected. Select Automatic to enable the controller port to automatically exchange information about data transmission speed and duplex capabilities. Auto negotiation is helpful when in an environment where different devices are connected and disconnected on a regular basis. Automatic is the default setting. |
| Duplex | Select either half, full or automatic as the duplex option. Select Half duplex to send data over the port, then immediately receive data from the same direction in which the data was transmitted. Like a full-duplex transmission, a half-duplex transmission can carry data in both directions, just not at the same time. Select Full duplex to transmit data to and from the controller port at the same time. Using full duplex, the port can send data while receiving data as well. Select Automatic to enable to the controller to dynamically duplex as port performance needs dictate. Automatic is the default setting. |

10. Enable or disable the following **CDP/LLDP** parameters used to configure Cisco Discovery Protocol and Link Layer Discovery Protocol for this profile's Ethernet port configuration:

- | | |
|---|---|
| Cisco Discovery Protocol Receive | Select this box to allow the Cisco discovery protocol to be received on this controller port. If enabled, the port sends out periodic interface updates to a multicast address to advertise its presence to neighbors. This option is enabled by default. |
|---|---|

Cisco Discovery Protocol Transmit	Select this box to allow the Cisco discovery protocol to be transmitted on this controller port. If enabled, the port sends out periodic interface updates to a multicast address to advertise its presence to neighbors.
Link Layer Discovery Protocol Receive	Select this box to allow the Link Layer discovery protocol to be received on this controller port. If enabled, the port sends out periodic interface updates to a multicast address to advertise its presence to neighbors. This option is enabled by default.
Link Layer Discovery Protocol Transmit	Select this box to allow the Link Layer discovery protocol to be transmitted on this controller port. If enabled, the port sends out periodic interface updates to a multicast address to advertise its presence to neighbors.

11. Set or override the following **Power Over Ethernet (PoE)** parameters used with this profile's Ethernet port configuration:

Enable POE	Select the check box to configure the selected port to use Power over Ethernet. To disable PoE on a port, uncheck this option. Power over Ethernet is supported on RFS4000 and RFS6000 model controllers only. When enabled, the controller supports 802.3af PoE on each of its ge ports. The PoE allows users to monitor port power consumption and configure power usage limits and priorities for each ge port.
Power Limit	Use the spinner control to set the total watts available for Power over Ethernet on the controller ge port. Set a value between 0 - 40 watts.
Power Priority	Set the power priority for the listed port to either to either <i>Critical</i> , <i>High</i> or <i>Low</i> . This is the priority assigned to this port versus the power requirements of the other supports available on the controller.

12. Define or override the following **Switching Mode** parameters applied to the Ethernet port configuration:

Mode	Select either the Access or Trunk radio button to set the VLAN switching mode over the port. If Access is selected, the port accepts packets only form the native VLANs. Frames are forwarded out the port untagged with no 802.1Q header. All frames received on the port are expected as untagged and are mapped to the native VLAN. If the mode is set to Trunk, the port allows packets from a list of VLANs you add to the trunk. A port configured as Trunk supports multiple 802.1Q tagged VLANs and one Native VLAN which can be tagged or untagged. Access is the default mode.
-------------	--

Native VLAN

Use the spinner control to define a numerical **Native VLAN ID** between 1 - 4094. The native VLAN allows an Ethernet device to associate untagged frames to a VLAN when no 802.1Q frame is included in the frame. Additionally, the native VLAN is the VLAN which untagged traffic will be directed over when using a port in trunk mode. The default VLAN is 1.

Tag Native VLAN

Select the check box to tag the native VLAN. Controllers support the IEEE 802.1Q specification for tagging frames and coordinating VLANs between devices. IEEE 802.1Q adds four bytes to each frame identifying the VLAN ID for upstream devices that the frame belongs. If the upstream Ethernet device does not support IEEE 802.1Q tagging, it does not interpret the tagged frames. When VLAN tagging is required between devices, both devices must support tagging and be configured to accept tagged VLANs. When a frame is tagged, the 12 bit frame VLAN ID is added to the 802.1Q header so upstream Ethernet devices know which VLAN ID the frame belongs to. The device reads the 12 bit VLAN ID and forwards the frame to the appropriate VLAN. When a frame is received with no 802.1Q header, the upstream device classifies the frame using the default or native VLAN assigned to the Trunk port. The native VLAN allows an Ethernet device to associate untagged frames to a VLAN when no 802.1Q frame is included in the frame. This feature is disabled by default.

Allowed VLANs

Selecting Trunk as the mode enables the **Allowed VLANs** parameter. Add VLANs that exclusively send packets over the listed port.

13. Optionally select the **Port Channel** checkbox from the **Port Channel Membership** area and define or override a setting between 1 - 8 using the spinner control. This sets the channel group for the port.
14. Select **OK** to save the changes and overrides made to the profile's Ethernet Port Basic Configuration. Select **Reset** to revert to the last saved configuration.
15. Select the **Security** tab.

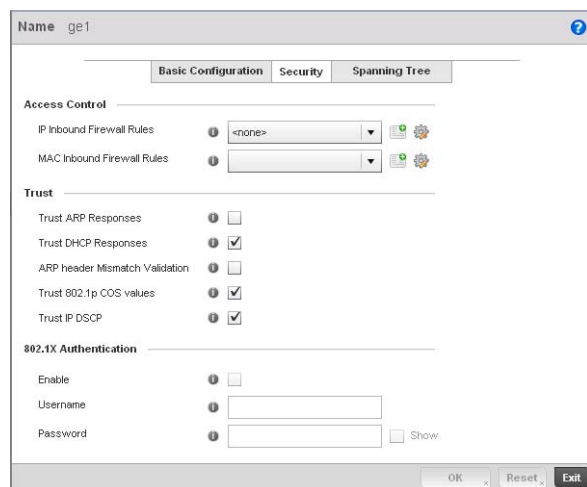


FIGURE 115 Profile Overrides - Ethernet Ports Security screen

16. Refer to the **Access Control** field. As part of the port's security configuration, Inbound IP and MAC address firewall rules are required.

Use the **Inbound IP Firewall Rules** and **Inbound MAC Firewall Rules** drop-down menus to select or override the firewall rules applied to this profile's Ethernet port configuration.

The firewall inspects IP and MAC traffic flows and detects attacks typically not visible to traditional wired firewall appliances.

If a firewall rule does not exist suiting the data protection needs of the target port configuration, select the **Create** icon to define a new rule configuration or the **Edit** icon to update or override an existing configuration. For more information, see *Wireless Firewall on page 9-419*.

17. Refer to the **Trust** field to define or override the following:

Trust ARP Responses	Select the check box to enable ARP trust on this controller port. ARP packets received on this port are considered trusted, and the information from these packets is used to identify rogue devices within the managed network. The default value is disabled.
Trust DHCP Responses	Select the check box to enable DHCP trust on this port. If enabled, only DHCP responses are trusted and forwarded on this port, and a DHCP server can be connected only to a DHCP trusted port. The default value is enabled.
ARP header Mismatch Validation	Select the check box to enable a mismatch check for the source MAC in both the ARP and Ethernet header. The default value is enabled.
Trust 8021p COS values	Select the check box to enable 802.1p COS values on this port. The default value is enabled.
Trust IP DSCP	Select the check box to enable IP DSCP values on this port. The default value is enabled.

NOTE

Some vendor solutions with VRRP enabled send ARP packets with Ethernet SMAC as a physical MAC and inner ARP SMAC as VRRP MAC. If this configuration is enabled, a packet is allowed, despite a conflict existing.

18. Set the following **802.1X Authentication** settings for the WLAN's QoS policy:

Enable	Check this box to enable 802.1X authentication on the selected Ethernet Port.
Username	Specify the configured username on the RADIUS server configured for authentication on this Ethernet port. This option is only available when the Enable checkbox is selected.
Password	Specify the password for the configured username used for authentication on this Ethernet port. This option is only available when the Enable checkbox is selected.

19. Select **OK** to save the changes and overrides made to the Ethernet port's security configuration. Select **Reset** to revert to the last saved configuration.

20. Select the **Spanning Tree** tab.

FIGURE 116 Profile Overrides - Ethernet Ports Spanning Tree screen

21. Set or override the following parameters for the port's **MSTP configuration**:

- Enable as Edge Port** Select the check box to define this port as an edge port. Using an edge (private) port, you can isolate devices to prevent connectivity over this port.
- Link Type** Select either the **Point-to-Point** or **Shared** radio button. Selecting Point-to-Point indicates the port should be treated as connected to a point-to-point link. Selecting Shared indicates this port should be treated as having a shared connection. A port connected to a hub is on a shared link, while one connected to a controller is a point-to-point link.
- Enable Cisco MSTP Interoperability** Select either the **Enable** or **Disable** radio buttons. This enables interoperability with Cisco's version of MSTP over the port, which is incompatible with standard MSTP.
- Force Protocol Version** Sets the protocol version to either *STP(0)*, *Not Supported(1)*, *RSTP(2)* or *MSTP(3)*. MSTP is the default setting.
- Guard** Determines whether the port enforces root bridge placement. Setting the guard to **Root** ensures the port is a designated port. Typically, each guard root port is a designated port, unless two or more ports (within the root bridge) are connected together. If the bridge receives superior (BPDUs) on a guard root-enabled port, the guard root moves the port to a root-inconsistent STP state. This state is equivalent to a listening state. No data is forwarded across the port. Thus, the guard root enforces the root bridge position.
- Enable PortFast** Select the check box to enable drop-down menus for both the Enable Portfast BPDU Filter and Enable Portfast BPDU guard options for the controller port.
- PortFast BPDU Filter** Select enable to invoke a BPDU filter for this portfast enabled port. Enabling the BPDU filter feature ensures this PortFast enabled port does not transmit or receive BPDUs.
- PortFast BPDU Guard** Select enable to invoke a BPDU guard for this portfast enabled port. Enabling the BPDU Guard feature means this portfast-enabled port will shutdown on receiving a BPDU.

22. Refer to the **Spanning Tree Port Cost** table.

Define or override an **Instance Index** using the spinner control and then set the **Cost**. The default path cost depends on the user defined speed of the port. The cost helps determine the role of the port in the MSTP network. The designated cost is the cost for a packet to travel from this port to the root in the MSTP configuration. The slower the media, the higher the cost.

Speed	Default Path Cost
<=100000 bits/sec	200000000
<=1000000 bits/sec	20000000
<=10000000 bits/sec	2000000
<=100000000 bits/sec	200000
<=1000000000 bits/sec	20000
<=10000000000 bits/sec	2000
<=100000000000 bits/sec	200
<=1000000000000 bits/sec	20
>1000000000000 bits/sec	2

23. Select **+ Add Row** as needed to include additional indexes.

24. Refer to the **Spanning Tree Port Priority** table.

Define or override an **Instance Index** using the spinner control and then set the **Priority**. The lower the priority, a greater likelihood of the port becoming a designated port. Thus applying an higher override value impacts the port's likelihood of becoming a designated port.

25. Select **+ Add Row** needed to include additional indexes.

26. Select **OK** to save the changes and overrides made to the Ethernet Port Spanning Tree configuration. Select **Reset** to revert to the last saved configuration.

Virtual Interface Override Configuration

Profile Interface Override Configuration

A controller Virtual Interface is required for layer 3 (IP) access to the controller or provide layer 3 service on a VLAN. The Virtual Interface defines which IP address is associated with each VLAN ID the controller is connected to. A Virtual Interface is created for the default VLAN (VLAN 1) to enable remote controller administration. A Virtual Interface is also used to map VLANs to IP address ranges. This mapping determines the destination for controller routing.

To review existing Virtual Interface configurations and either create a new Virtual Interface configuration, modify (override) an existing configuration or delete an existing configuration:

1. Select the **Configuration** tab from the Web UI.
2. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

3. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

Once the configurations of existing Virtual Interfaces have been reviewed, determine whether a new interface requires creation, or an existing Virtual Interface requires edit (override) or deletion.

8. Select **Add** to define a new Virtual Interface configuration, **Edit** to modify or override the configuration of an existing Virtual Interface or **Delete** to permanently remove a selected Virtual Interface.

FIGURE 118 Profile Overrides - Virtual Interfaces Basic Configuration screen

9. The **Basic Configuration** screen displays by default regardless of whether a new Virtual Interface is being created or an existing one is being modified.
10. If creating a new Virtual Interface, use the **VLAN ID** spinner control to define a numeric VLAN ID between 1 - 4094.
11. Define or override the following parameters from within the **Properties** field:

Description Provide or edit a description (up to 64 characters) for the Virtual Interface that helps differentiate it from others with similar configurations.

Admin Status Either select the *Disabled* or *Enabled* radio button to define this interface's current status within the managed network. When set to Enabled, the Virtual Interface is operational and available to the controller. The default value is disabled.

12. Set or override the following network information from within the **IP Addresses** field:

Enable Zero Configuration Define the IP address for the VLAN associated Virtual Interface.

Primary IP Address Define the IP address for the VLAN associated Virtual Interface.

Use DHCP to Obtain IP Select this option to allow DHCP to provide the IP address for the Virtual Interface. Selecting this option disables the Primary IP address field.

Use DHCP to obtain Gateway/DNS Servers Select this option to allow DHCP to obtain a default gateway address, and DNS resource for one virtual interface. This setting is disabled by default and only available when the Use DHCP to Obtain IP option is selected.

Secondary Addresses Use the Secondary Addresses parameter to define additional IP addresses to associate with VLAN IDs. The address provided in this field is used if the primary IP address is unreachable.

13. Refer to the **DHCP Relay** field to set or override the DHCP relay server configuration used with the controller Virtual Interface.

Respond to DHCP Relay Packets Select the Respond to DHCP Relay Packets option to allow the controller's onboard DHCP server to respond to relayed DHCP packets on this interface.

DHCP Relay IP Address Provide IP addresses for DHCP server relay resources.
The interface VLAN and gateway should have their IP addresses set. The interface VLAN and gateway interface should not have DHCP client or DHCP Server enabled. DHCP packets cannot be relayed to an onboard DHCP Server. The interface VLAN and gateway interface cannot be the same.
When changing from a default DHCP address to a fixed IP address, set a static route first. This is critical when the controller is being accessed from a subnet not directly connected to the controller and the default route was set from DHCP.

14. Define or override the **Network Address Translation (NAT)** direction.

Select either the **Inside**, **Outside** or **None** radio buttons.

- *Inside* - The inside network is transmitting data over the network its intended destination. On the way out, the source IP address is changed in the header and replaced by the (public) IP address.
- *Outside* - Packets passing through the NAT on the way back to the managed LAN are searched against to the records kept by the NAT engine. There the destination IP address is changed back to the specific internal private class IP address in order to reach the LAN over the switch managed network.
- *None* - No NAT activity takes place. This is the default setting.

NOTE

Refer to [Setting the Profile's NAT Configuration on page 7-383](#) for instructions on creating a profile's NAT configuration.

15. Select the **OK** button to save the changes and overrides to the Basic Configuration screen.
Select **Reset** to revert to the last saved configuration.

16. Select the **Security** tab.



FIGURE 119 Profile Overrides - Virtual Interfaces Security screen

17. Use the **Inbound IP Firewall Rules** drop-down menu to select the firewall rule configuration to apply to this Virtual Interface.

The firewall inspects and packet traffic to and from connected clients.

If a firewall rule does not exist suiting the data protection needs of this Virtual Interface, select the **Create** icon to define a new firewall rule configuration or the **Edit** icon to modify or override an existing configuration. For more information, see [Wireless Firewall on page 9-419](#).

18. Use the **VPN Crypto Map** drop-down menu to select or override the Crypto Map configuration applied to this Virtual Interface.

Crypto Map entries are sets of configuration parameters for encrypting packets that pass through the VPN Tunnel. If a Crypto Map configuration does not exist suiting the needs of this Virtual Interface, select the **Create** icon to define a new Crypto Map configuration or the **Edit** icon to modify an existing configuration. For more information, see [Overriding a Profile's VPN Configuration on page 5-189](#).

19. Select the **OK** button located at the bottom right of the screen to save the changes and overrides to the Security screen. Select **Reset** to revert to the last saved configuration.

Port Channel Override Configuration

Profile Interface Override Configuration

Controller profiles can be customized port channel configurations as part of their interface configuration. Existing port channel profile configurations can be overridden as they become obsolete for specific device deployments.

To define or override a port channel configuration on a controller profile:

1. Select the **Configuration** tab from the Web UI.
2. Select **Devices** from the Configuration tab.


The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

3. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

4. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
5. Select **Interface** to expand its sub menu options.
6. Select **Port Channels**.

The Port Channels screen displays.

Name	Type	Description	Admin Status
port-channel1	Port Channel	test1	 Enabled

Type to search in tables

Row Count: 1

Add Edit Delete Exit

FIGURE 120 Profile Overrides - Port Channels screen

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

7. Refer to the following to review existing port channel configurations and status to determine whether a parameter requires an override:

- | | |
|---------------------|---|
| Name | Displays the port channel’s numerical identifier assigned when it was created. The numerical name cannot be modified as part of the edit process. |
| Type | Displays whether the type is port channel. |
| Description | Lists a a short description (64 characters maximum) describing the port channel or differentiating it from others with similar configurations. |
| Admin Status | A green checkmark defines the listed port channel as active and currently enabled with the controller profile. A red “X” defines the port channel as currently disabled and not available for use. The interface status can be modified with the port channel configuration as required |

8. To edit or override the configuration of an existing port channel, select it from amongst those displayed and select the **Edit** button. The port channel **Basic Configuration** screen displays by default.

FIGURE 121 Profile Overrides - Port Channels Basic Configuration screen

9. Set or override the following port channel **Properties**:

- Description** Enter a brief description for the controller port channel (64 characters maximum).
- Admin Status** Select the **Enabled** radio button to define this port channel as active to the controller profile it supports. Select the **Disabled** radio button to disable this port channel configuration in the controller profile. It can be activated at any future time when needed. The default setting is disabled.
- Speed** Select the speed at which the port channel can receive and transmit data. Select either 10 Mbps, 100 Mbps or 1000 Mbps to establish a 10, 100 or 1000 Mbps data transfer rate for the selected half duplex or full duplex transmission. These options are not available if Auto is selected. Select Automatic to allow the port channel to automatically exchange information about data transmission speeds and duplex capabilities. Auto negotiation is helpful in an environment where different devices are connected and disconnected on a regular basis. Automatic is the default setting.
- Duplex** Select either half, full or automatic as the duplex option. Select **Half** duplex to send data over the port channel, then immediately receive data from the same direction in which the data was transmitted. Like a full-duplex transmission, a half-duplex transmission can carry data in both directions, just not at the same time. Select **Full** duplex to transmit data to and from the port channel at the same time. Using full duplex, the port channel can send data while receiving data as well. Select **Automatic** to enable to the controller to dynamically duplex as port channel performance needs dictate. Automatic is the default setting.

10. Use the **Port Channel Load Balance** drop-down menu from the **Client Load Balancing** section to define whether port channel load balancing is conducted using a *Source/Destination IP* or a *Source/Destination MAC*. Source/Destination IP is the default setting.

11. Define or override the following **Switching Mode** parameters to apply to the port channel configuration:

Mode	Select either the Access or Trunk radio button to set the VLAN switching mode over the port channel. If Access is selected, the port channel accepts packets only from the native VLAN. Frames are forwarded untagged with no 802.1Q header. All frames received on the port are expected as untagged and are mapped to the native VLAN. If the mode is set to Trunk, the port channel allows packets from a list of VLANs you add to the trunk. A port channel configured as Trunk supports multiple 802.1Q tagged VLANs and one Native VLAN which can be tagged or untagged. Access is the default setting.
Native VLAN	Use the spinner control to define a numerical Native VLAN ID between 1 - 4094. The native VLAN allows an Ethernet device to associate untagged frames to a VLAN when no 802.1Q frame is included in the frame. Additionally, the native VLAN is the VLAN untagged traffic will be directed over when using trunk mode. The default value is 1.
Tag the Native VLAN	Select the checkbox to tag the native VLAN. Controllers support the IEEE 802.1Q specification for tagging frames and coordinating VLANs between devices. IEEE 802.1Q adds four bytes to each frame identifying the VLAN ID for upstream devices that the frame belongs. If the upstream Ethernet device does not support IEEE 802.1Q tagging, it does not interpret the tagged frames. When VLAN tagging is required between devices, both devices must support tagging and be configured to accept tagged VLANs. When a frame is tagged, a 12 bit frame VLAN ID is added to the 802.1Q header, so upstream Ethernet devices know which VLAN ID the frame belongs to. The device reads the 12 bit VLAN ID and forwards the frame to the appropriate VLAN. When a frame is received with no 802.1Q header, the upstream device classifies the frame using the default or native VLAN assigned to the Trunk port. The native VLAN allows an Ethernet device to associate untagged frames to a VLAN when no 802.1Q frame is included in the frame. This setting is disabled by default.
Allowed VLANs	Selecting Trunk as the mode enables the Allowed VLANs parameter. Add VLANs that exclusively send packets over the port channel.

12. Select **OK** to save the changes and overrides to the port channel Basic Configuration. Select **Reset** to revert to the last saved configuration.

13. Select the **Security** tab.

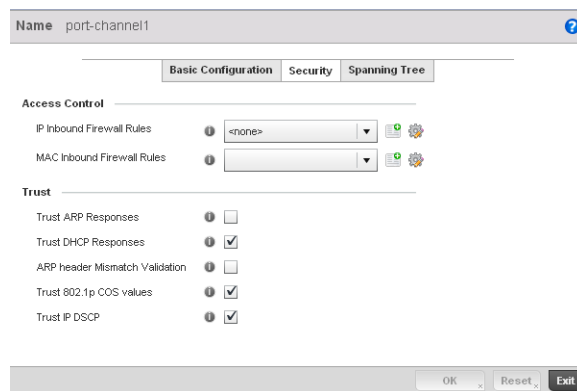


FIGURE 122 Profile Overrides - Port Channels Security screen

14. Refer to the **Access Control** field. As part of the port channel's security configuration, Inbound IP and MAC address firewall rules are required.

Use the **Inbound IP Firewall Rules** and **Inbound MAC Firewall Rules** drop-down menus to select or override the firewall rules to apply to this profile's port channel configuration.

The firewall inspects IP and MAC traffic flows and detects attacks typically not visible to traditional wired firewall appliances.

If a firewall rule does not exist suiting the data protection needs of the target port channel configuration, select the **Create** icon to define a new rule configuration, or the **Edit** icon to modify (override) an existing firewall rule configuration. For more information, see [Wireless Firewall on page 9-419](#).

15. Refer to the **Trust** section to define or override the following:

- Trust ARP Responses** Select the check box to enable ARP trust on this port channel. ARP packets received on this controller port are considered trusted, and information from these packets is used to identify rogue devices within the managed network. The default value is disabled.
- Trust DHCP Responses** Select the check box to enable DHCP trust. If enabled, only DHCP responses are trusted and forwarded on this port channel, and a DHCP server can be connected only to a DHCP trusted port. The default value is enabled.
- ARP header Mismatch Validation** Select the check box to enable a mismatch check for the source MAC in both the ARP and Ethernet header. The default value is enabled.
- Trust 802.1p COS values** Select the check box to enable 802.1p COS values on this port channel. The default value is enabled.
- Trust IP DSCP** Select the check box to enable IP DSCP values on this port channel. The default value is disabled.

16. Select **OK** to save the changes and overrides to the security configuration. Select **Reset** to revert to the last saved configuration.

17. Select the **Spanning Tree** tab.

FIGURE 123 Profile Overrides - Port Channels Spanning Tree screen

18. Define or override the following **PortFast** parameters for the port channel's MSTP configuration:

Enable PortFast	Select the check box to enable drop-down menus for both the Enable Portfast BPDU Filter and Enable Portfast BPDU guard options for the controller port. This setting is disabled by default.
PortFast BPDU Filter	Select enable to invoke a BPDU filter for this portfast enabled port channel. Enabling the BPDU filter feature ensures this port channel does not transmit or receive any BPDUs. The default setting is None.
PortFast BPDU Guard	Select enable to invoke a BPDU guard for this portfast enabled port channel. Enabling the BPDU Guard feature means this port will shutdown on receiving a BPDUs. Hence no BPDUs are processed. The default setting is None.

19. Set or override the following **MSTP Configuration** parameters for the port channel:

Enable as Edge Port	Select the check box to define this port as an edge port. Using an edge (private) port, you can isolate devices to prevent connectivity over this port channel. This setting is disabled by default.
Link Type	Select either the Point-to-Point or Shared radio button. Selecting Point-to-Point indicates the port should be treated as connected to a point-to-point link. Selecting Shared indicates this port should be treated as having a shared connection. A port connected to a hub is on a shared link, while one connected to a controller is a point-to-point link. Point-to-Point is the default setting.
Cisco MSTP Interoperability	Select either the Enable or Disable radio buttons. This enables interoperability with Cisco's version of MSTP, which is incompatible with standard MSTP. This setting is disabled by default.
Force Protocol Version	Sets the protocol version to either STP(0), Not Supported(1), RSTP(2) or MSTP(3). MSTP is the default setting.
Guard	Determines whether the port channel enforces root bridge placement. Setting the guard to Root ensures the port is a designated port. Typically, each guard root port is a designated port, unless two or more ports (within the root bridge) are connected together. If the bridge receives superior (BPDUs) on a guard root-enabled port, the guard root moves the port to a root-inconsistent STP state. This state is equivalent to a listening state. No data is forwarded across the port. Thus, the guard root enforces the root bridge position.

20. Refer to the **Spanning Tree Port Cost** table.

Define or override an **Instance Index** using the spinner control and then set the **Cost**. The default path cost depends on the user defined port speed. The cost helps determine the role of the port channel in the MSTP network. The designated cost is the cost for a packet to travel from this port to the root in the MSTP configuration. The slower the media, the higher the cost.

Speed	Default Path Cost
<=100000 bits/sec	200000000
<=1000000 bits/sec	20000000
<=10000000 bits/sec	2000000
<=100000000 bits/sec	200000
<=1000000000 bits/sec	20000

Speed	Default Path Cost
<=100000000000 bits/sec	2000
<=100000000000 bits/sec	200
<=1000000000000 bits/sec	20
>1000000000000 bits/sec	2

Select **+ Add Row** as needed to include additional indexes.

Refer to the **Spanning Tree Port Priority** table.

Define or override an **Instance Index** using the spinner control and then set the **Priority**. The lower the priority, a greater likelihood of the port becoming a designated port.

21. Select **+ Add Row** as needed to include additional indexes.
22. Select **OK** to save the changes and overrides made to the Ethernet Port Spanning Tree configuration. Select **Reset** to revert to the last saved configuration.

Radio Override Configuration

Profile Interface Override Configuration

Access Points can have their radio profile configurations overridden once their radios have successfully associated to the managed network.

To define a radio configuration override from the Access Point's associated controller:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.
2. Select a target Access Point (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.
3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Interface** to expand its sub menu options.
5. Select **Radios**.





Name 	Type	Description	Admin Status	RF Mode	Channel	Transmit Power
radio1	Radio	radio1	 Disabled		random-width-40	20
radio2	Radio		 Enabled		smart	20
radio3	Radio		 Enabled		smart	20
Type to search in tables		Row Count: 3				

FIGURE 124 Profile Overrides - Access Point Radios screen

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

- Review the following radio configuration data to determine whether a radio configuration requires modification or override to better support the managed network:

Name	Displays whether the reporting radio is the Access Point's radio1, radio2 or radio3. Legacy AP-7131 models contain either a single or a dual radio configuration. Newer AP-7131N model Access Points support single, dual or triple radio configurations. An Brocade Mobility 650 Access Point model Access Point is available in either single or dual radio models.
Type	Displays the type of radio housed by each listed Access Point.
Description	Displays a brief description of the radio provided by the administrator when the radio's configuration was added or modified.
Admin Status	A green checkmark defines the listed Virtual Interface configuration as active and enabled with its supported controller profile. A red "X" defines the Virtual Interface as currently disabled. The interface status can be modified when a new Virtual Interface is created or an existing one modified.

- RF Mode** Displays whether each listed radio is operating in the 802.11a/n or 802.11b/g/n radio band. If the radio is a dedicated sensor, it will be listed as a sensor to define the radio as not providing typical WLAN support. The radio band is set from within the Radio Settings tab.
- Channel** Lists the channel setting for the radio. Smart is the default setting. If set to smart, the Access Point scans non-overlapping channels listening for beacons from other Access Points. After the channels are scanned, it selects the channel with the fewest Access Points. In the case of multiple access points on the same channel, it will select the channel with the lowest average power level. The column displays smart if set for dynamic Smart RF support.
- Transmit Power** Lists the transmit power for each radio displayed as a value in milliwatts.

7. If required, select a radio configuration and select the **Edit** button to modify or override portions of its configuration.

FIGURE 125 Profile Overrides - Access Point Radio Settings tab

The **Radio Settings** tab displays by default.

8. Define or override the following radio configuration parameters from within the **Properties** field:

- Description** Provide or edit a description (1 - 64 characters in length) for the radio that helps differentiate it from others with similar configurations.

Admin Status	Either select the Active or Shutdown radio button to define this radio's current status within the managed network. When defined as Active, the Access Point is operational and available for client support within the managed network.
Radio QoS Policy	Use the drop-down menu to specify an existing QoS policy to apply to the Access Point radio in respect to its intended radio traffic. If there's no existing suiting the radio's intended operation, select the Create icon to define a new QoS policy that can be applied to this controller profile. For more information, see Radio QoS Policy on page 6-278 .
Association ACL	Use the drop-down menu to specify an existing Association ACL policy to apply to the Access Point radio. An Association ACL is a policy-based <i>Access Control List</i> (ACL) that either prevents or allows wireless clients from connecting to a managed Access Point radio. An ACL is a sequential collection of permit and deny conditions that apply to controller packets. When a packet is received on an interface, the controller compares the fields in the packet against any applied ACLs to verify the packet has the required permissions to be forwarded, based on the criteria specified in the access lists. If a packet does not meet any of the criteria specified in the ACL, the packet is dropped. Select the Create icon to define a new Association ACL that can be applied to this controller profile. For more information, see Association ACL on page 6-298 .
9. Set or override the following profile Radio Settings for the selected Access Point radio.	
RF Mode	Set the mode to either 2.4 GHz WLAN or 5 GHz WLAN support depending on the radio's intended client support. Set the mode to Sensor if using the radio for rogue device detection. To a radio as a detector, disable Sensor support on the other Access Point radio.
Lock RF Mode	Select the check box to lock Smart RF for this radio. The default setting is disabled.
Channel	Use the drop-down menu to select the channel of operation for the radio. Only a trained installation professional should define the radio channel. Select Smart for the radio to scan non-overlapping channels listening for beacons from other Access Points. After channels are scanned, the radio selects the channel with the fewest Access Points. In the case of multiple Access Points on the same channel, it selects the channel with the lowest average power level. The default value is Smart.
Transmit Power	Set the transmit power of the selected Access Point radio. If using a dual or three radio model Access Point, each radio should be configured with a unique transmit power in respect to its intended client support function. If using Smart RF select the Smart RF check box to let Smart RF determine the transmit power. A setting of 0 defines the radio as using Smart RF to determine its output power. 20 dBm is the default value.
Antenna Gain	Set the antenna between 0.00 - 15.00 dBm. The access point's <i>Power Management Antenna Configuration File</i> (PMACF) automatically configures the access point's radio transmit power based on the antenna type, its antenna gain (provided here) and the deployed country's regulatory domain restrictions. Once provided, the access point calculates the power range. Antenna gain relates the intensity of an antenna in a given direction to the intensity that would be produced ideally by an antenna that radiates equally in all directions (isotropically), and has no losses. Although the gain of an antenna is directly related to its directivity, its gain is a measure that takes into account the efficiency of the antenna as well as its directional capabilities. Brocade Mobility recommends that only a professional installer set the antenna gain. The default value is 0.00.
Antenna Mode	Set the number of transmit and receive antennas on the Access Point. 1x1 is used for transmissions over just the single "A" antenna, 1x3 is used for transmissions over the "A" antenna and all three antennas for receiving. 2x2 is used for transmissions and receipts over two antennas for dual antenna models. The default setting is dynamic based on the Access Point model deployed and its transmit power settings.

Enable Antenna Diversity	Select this box to enable antenna diversity on supported antennas. Antenna diversity uses two or more antennas to increase signal quality and strength. This option is disabled by default.
Wireless Client Power	Select this box to specify the transmit power on supported wireless clients. If this is enabled set a client power level between 0 to 20 dBm. This option is disabled by default.
Dynamic Chain Selection	Select this box for the radio to dynamically change the number of transmit chains. This option is enabled by default.
Rate	Use the Select button to set rate options depending on the 802.11 protocols selected. If the radio band is set to Sensor or Detector, the Data Rates drop-down menu is not enabled, as the rates are fixed and not user configurable. If 2.4 GHz is selected as the radio band, select separate 802.11b, 802.11g and 802.11n rates and define how they are used in combination. If 5 GHz is selected as the radio band, select separate 802.11a and 802.11n rates then define how they are used together. When using 802.11n (in either the 2.4 or 5 GHz band), Set a MCS (modulation and coding scheme) in respect to the radio's channel width and guard interval. A MCS defines (based on RF channel conditions) an optimal combination of 8 data rates, bonded channels, multiple spatial streams, different guard intervals and modulation types. Clients can associate as long as they support basic MCS (as well as non-11n basic rates).
Radio Placement	Use the drop-down menu to specify whether the radio is located Indoors or Outdoors. The placement should depend on the country of operation selected and its regulatory domain requirements for radio emissions. The default setting is Indoors.
Max Clients	Use the spinner control to set a maximum permissible number of clients to connect with this radio. The available range is between 0 - 256 clients. The default value is 256.

10. Set or override the following profile **WLAN Properties** for the selected Access Point radio:

Beacon Interval	Set the interval between radio beacons in milliseconds (either 50, 100 or 200). A beacon is a packet broadcast by adopted radios to keep the network synchronized. Included in a beacon is information such as the WLAN service area, the radio address, the broadcast destination addresses, a time stamp, and indicators about traffic and delivery (such as a DTIM). Increase the DTIM/beacon settings (lengthening the time) to let nodes sleep longer and preserve battery life. Decrease these settings (shortening the time) to support streaming-multicast audio and video applications that are jitter-sensitive. The default value is 100 milliseconds.
DTIM Interval BSSID	Set a DTIM Interval to specify a period for <i>Delivery Traffic Indication Messages</i> (DTIM). A DTIM is periodically included in a beacon frame transmitted from adopted radios. The DTIM indicates broadcast and multicast frames (buffered at the Access Point) are soon to arrive. These are simple data frames that require no acknowledgement, so nodes sometimes miss them. Increase the DTIM/ beacon settings (lengthening the time) to let nodes sleep longer and preserve their battery life. Decrease these settings (shortening the time) to support streaming multicast audio and video applications that are jitter-sensitive.
RTS Threshold	<p>Specify a <i>Request To Send</i> (RTS) threshold (between 1 - 2,347 bytes) for use by the WLAN's adopted Access Point radios. RTS is a transmitting station's signal that requests a Clear To Send (CTS) response from a receiving client. This RTS/CTS procedure clears the air where clients are contending for transmission time. Benefits include fewer data collisions and better communication with nodes that are hard to find (or hidden) because of other active nodes in the transmission path.</p> <p>Control RTS/CTS by setting an RTS threshold. This setting initiates an RTS/CTS exchange for data frames larger than the threshold, and sends (without RTS/CTS) any data frames smaller than the threshold.</p> <p>Consider the trade-offs when setting an appropriate RTS threshold for the WLAN's Access Point radios. A lower RTS threshold causes more frequent RTS/CTS exchanges. This consumes more bandwidth because of additional latency (RTS/CTS exchanges) before transmissions can commence. A disadvantage is the reduction in data-frame throughput. An advantage is quicker system recovery from electromagnetic interference and data collisions. Environments with more wireless traffic and contention for transmission make the best use of a lower RTS threshold.</p> <p>A higher RTS threshold minimizes RTS/CTS exchanges, consuming less bandwidth for data transmissions. A disadvantage is less help to nodes that encounter interference and collisions. An advantage is faster data-frame throughput. Environments with less wireless traffic and contention for transmission make the best use of a higher RTS threshold.</p>
Short Preamble	If using an 802.11bg radio, select this checkbox for the radio to transmit using a short preamble. Short preambles improve throughput. However, some devices (SpectraLink phones) require long preambles. The default value is disabled.
Guard Interval	Use the drop-down menu to specify a <i>Long</i> or <i>Any</i> guard interval. The guard interval is the space between characters being transmitted. The guard interval eliminates <i>inter-symbol interference</i> (ISI). ISI occurs when echoes or reflections from one character interfere with another character. Adding time between transmissions allows echo's and reflections to settle before the next character is transmitted. A shorter guard interval results in shorter character times which reduces overhead and increases data rates by up to 10%. The default value is Long.
Probe Response Rate	Use the drop-down menu to specify the data transmission rate used for the transmission of probe responses. Options include, highest-basic, lowest-basic and follow-probe-request (default setting).
Probe Response Retry	Select the check box to retry probe responses if they are not acknowledged by the target wireless client. The default value is enabled.

11. Select the **Enable Off Channel Scan** check box in the **Channel Scanning** section to enable scanning across all channels using this radio. Channel scans use Access Point resources and can be time consuming, so only enable when your sure the radio can afford the bandwidth be directed towards to the channel scan and does not negatively impact client support.
12. Select a mode from the **Feed WLAN Packets to Sensor** check box in the **Radio Share** section to enable this feature. Select either **Inline** or **Promiscuous** mode to allow the packets the radio is switching to also be used by the WIPS analysis module. This feature can be enabled in two modes: an inline mode where the wips sensor receives the packets from the radios with radio operating in normal mode. A promiscuous mode where the radio is configured to a mode where it receives all packets on the channel whether the destination adres is the radio or not, and the wips module can analyze them.
13. Select the **WLAN Mapping** tab.

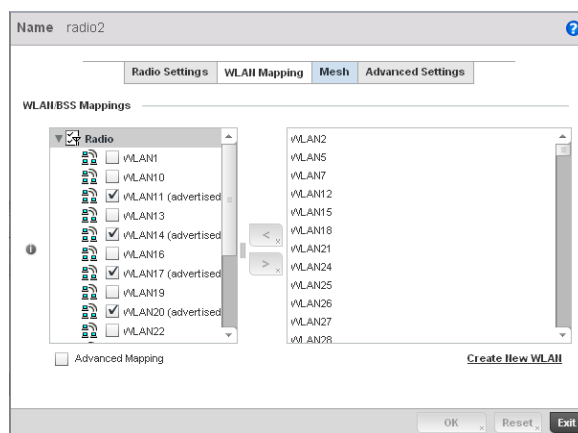


FIGURE 126 Profile Overrides - Access Point Radio WLAN Mapping tab

14. Refer to the **WLAN/BSS Mappings** field to set or override WLAN BSSID assignments for an existing Access Point deployment.

Administrators can assign each WLAN its own BSSID. If using a single-radio access point, there are 8 BSSIDs available. If using a dual-radio access point there are 8 BSSIDs for the 802.11b/g/n radio and 8 BSSIDs for the 802.11a/n radio.
15. Select **Advanced Mapping** to enable WLAN mapping to a specific BSS ID.
16. Select **OK** to save the changes and overrides to the WLAN Mapping. Select **Reset** to revert to the last saved configuration.
17. Select the **MeshConnex** tab.

FIGURE 127 Profile Overrides - Access Point Radio Mesh tab

18. Refer to the **Settings** field to define or override basic mesh settings for the Access Point radio.

Mesh

Use the pulldown to set the mesh mode for this radio. Available options are Disabled, Portal or Client. Setting the mesh mode to Disabled deactivates all mesh activity on this radio. Setting the mesh mode to Portal turns the radio into a mesh portal. This will start the radio beaconing immediately and will accept connections from other mesh nodes. Setting the mesh mode to client enables the radio to operate as a mesh client that will scan for and connect to mesh portals or nodes that are connected to portals.

Mesh Links

Specify the number of mesh links allowed by the radio. The radio can have between 1-6 mesh links when the radio is configured as a Portal.

NOTE

Only single hop mesh links are supported at this time.

NOTE

The mesh encryption key is configurable from the *Command Line Interface (CLI)* using the command 'mesh psk'. Administrators must ensure that this key is configured on the AP when it is being staged for mesh, and also added to the mesh client as well as to the portal APs configuration on the controller. For more information about the CLI please see the *Brocade Mobility RFS4000, RFS6000 and RFS7000 CLI Reference Guide*.

19. Refer to the **Preferred Peer Device** table to add mesh peers. For each peer being added enter its MAC Address and a Priority between 1 and 6. The lower the priority number the higher priority it'll be given when connecting to mesh infrastructure.

20. Select the **+ Add Row** button to add preferred peer devices for the radio to connect to in mesh mode.

21. Select the **Advanced Settings** tab.

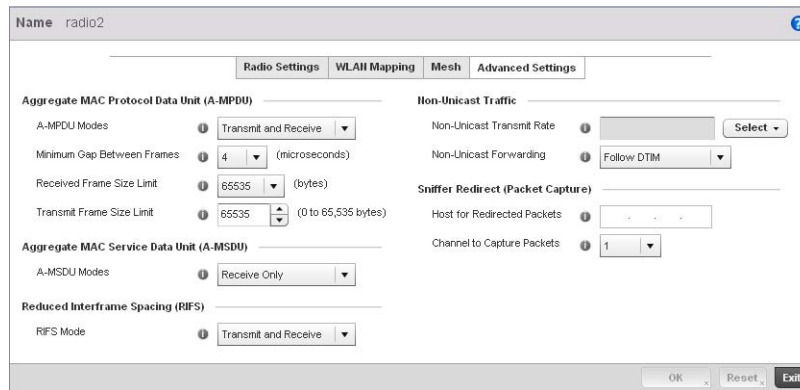


FIGURE 128 Profile Overrides - Access Point Radio Advanced Settings tab

22. Refer to the **Aggregate MAC Protocol Data Unit (A-MPDU)** field to define or override how MAC service frames are aggregated by the Access Point radio.

A-MPDU Modes

Use the drop-down menu to define the A-MPDU mode supported. Options include *Transmit Only*, *Receive Only*, *Transmit and Receive* and *None*. The default value is *Transmit and Receive*. Using the default value, long frames can be both sent and received (up to 64 KB). When enabled, define either a transmit or receive limit (or both).

Minimum Gap Between Frames

Use the drop-down menu to define the minimum gap between A-MPDU frames (in microseconds). The default value is 4 microseconds.

Received Frame Size Limit

If a support mode is enabled allowing A-MPDU frames to be received, define an advertised maximum limit for received A-MPDU aggregated frames. Options include 8191, 16383, 32767 or 65535 bytes. The default value is 65535 bytes.

Transmit Frame Size Limit

Use the spinner control to set limit on transmitted A-MPDU aggregated frames. The available range is between 0 - 65,535 bytes). The default value is 65535 bytes.

23. Use the **A-MSDU Modes** drop-down menu in the **Aggregate MAC Service Data Unit (A-MSDU)** section to set or override the supported A-MSDU mode.

Available modes include *Receive Only* and *Transmit and Receive*. *Transmit and Receive* is the default value. Using *Transmit and Receive*, frames up to 4 KB can be sent and received. The buffer limit is not configurable.

24. Define a **RIFS Mode** using the drop-down menu in the **Reduced Interframe Spacing (RIFS)** section. This value determines whether interframe spacing is applied to Access Point transmissions or received packets, or both or none. The default mode is *Transmit and Receive*.

Consider setting this value to *None* for high priority traffic to reduce packet delay.

25. Set or override the following **Non-Unicast Traffic** values for the profile's supported Access Point radio and its connected wireless clients:

Non-Unicast Transmit Rate

Use the **Select** drop-down menu to launch a sub screen to define the data rate for broadcast and multicast frame transmissions. Seven different rates are available if the not using the same rate for each BSSID, each with a separate menu.

Non-Unicast Forwarding

Define whether client broadcast and multicast packets should always follow DTIM, or only follow DTIM when using Power Save Aware mode. The default setting is *Follow DTIM*.

26. Refer to the **Sniffer Redirect (Packet Capture)** field to define or override the radio's captured packet configuration.

Host for Redirected Packets	If packets are re-directed from a controller's connected Access Point radio, define an IP address of a resource (additional host system) used to capture the re-directed packets. This address is the numerical (non DNS) address of the host used to capture the re-directed packets.
Channel to Capture Packets	Use the drop-down menu to specify the channel used to capture re-directed packets. The default value is channel 1.

27. Select **OK** to save or override the changes to the Advanced Settings screen. Select **Reset** to revert to the last saved configuration.

WAN Backhaul Override Configuration

Profile Interface Override Configuration

A *Wireless Wide Area Network (WWAN)* card is a specialized network interface card that allows a network device to connect, transmit and receive data over a Cellular Wide Area Network. The AP71xx, RFS4000 and RFS6000 all have a PCI Express card slot that supports 3G WWAN cards. The WWAN card uses point to point protocol (PPP) to connect to the Internet Service Provider (ISP) and gain access to the Internet. PPP is the protocol used for establishing internet links over dial-up modems, DSL connections, and many other types of point-to-point communications. PPP packages your system's TCP/IP packets and forwards them to the serial device where they can be put on the network. PPP is a full-duplex protocol that can be used on various physical media, including twisted pair or fiber optic lines or satellite transmission. It uses a variation of High Speed Data Link Control (HDLC) for packet encapsulation.

To define a WAN Backhaul configuration override:

1. Select **Devices** from the Configuration tab.
The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.
2. Select a target Access Point (by double-clicking it) from amongst those displayed within the Device Configuration screen.
Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.
3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Interface** to expand its sub menu options.
5. Select **WAN Backhaul**.

FIGURE 129 Profile Overrides -WAN Backhaul screen

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

6. Refer to the **WAN (3G) Backhaul** configuration to specify WAN card settings:

WAN Interface Name	Displays the WAN Interface name for the WAN 3G Backhaul card.
Reset WAN Card	If the WAN Card becomes unresponsive or is experiencing other errors click the Reset WAN Card button to power cycle and reboot the WAN card.
Enable WAN (3G)	Check this box to enable 3G WAN card support on the device. A supported 3G card must be connected to the device for this feature to work.

7. Define or override the following authentication parameters from within the **Basic Settings** field:

Username	Provide your username for authentication support by your cellular data carrier.
Password	Provide your password for authentication support by your cellular data carrier.
Access Point Name (APN)	Enter the name of the cellular data provider if necessary. This setting is needed in areas with multiple cellular data providers using the same protocols such as Europe, the middle east and Asia.
Authentication Type	Use the pull-down menu to specify authentication type used by your cellular data provider. Supported authentication types are None, PAP, CHAP, MSCHAP, and MSCHAP-v2.

8. Select **OK** to save or override the changes to the Advanced Settings screen. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's Network Configuration

Setting a profile's network configuration is a large task comprised of numerous controller administration activities. Each of the configuration activities described below can have an override applied to the original profile configuration. Applying an override removes the device from the profile configuration that may be shared by other devices and requires careful administration to ensure this one device still supports the deployment requirements within the managed network.

A profile's network configuration process consists of the following:

- [Overriding a Profile's DNS Configuration](#)
- [Overriding a Profile's ARP Configuration](#)
- [Overriding a Profile's Quality of Service \(QoS\) Configuration](#)
- [Overriding a Profile's Spanning Tree Configuration](#)
- [Overriding a Profile's Static Route Configuration](#)
- [Overriding a Profile's Forwarding Database Configuration](#)
- [Overriding a Profile's Bridge VLAN Configuration](#)
- [Overriding a Profile's Miscellaneous Network Configuration](#)

Overriding a Profile's DNS Configuration

[Overriding a Profile's Network Configuration](#)

Domain Naming System (DNS) DNS is a hierarchical naming system for resources connected to the Internet or a private network. Primarily, the controller's DNS resources translate domain names into IP addresses. If one DNS server doesn't know how to translate a particular domain name, it asks another one until the correct IP address is returned. DNS enables access to resources using human friendly notations. DNS converts human friendly domain names into notations used by different networking equipment for locating resources.

As a resource is accessed (using human-friendly hostnames), it's possible to access the resource even if the underlying machine friendly notation name changes. Without DNS you need to remember a series of numbers (123.123.123.123) instead of a domain name (www.domainname.com).

The controller maintains its own DNS facility that can assist in domain name translation. A DNS assignment can be overridden as needed, but removes the device configuration from the managed profile that may be shared with other similar device models.

To define the controller's DNS configuration or apply overrides to an existing configuration:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.
2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.
3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Network** to expand its sub menu options.

5. Select DNS.

FIGURE 130 Profile Overrides - Network DNS screen

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

6. Set or override the following controller **Domain Name System (DNS)** configuration data:

- Domain Name** Provide or override the default Domain Name used to resolve DNS names. The name cannot exceed 64 characters.
- Enable Domain Lookup** Select the check box to enable DNS on the controller. When enabled, the controller can convert human friendly domain names into numerical IP destination addresses. The check box is selected by default.
- DNS Server Forwarding** Click to enable the forwarding DNS queries to external DNS servers if a DNS query cannot be processed by the controller's own DNS resources. This feature is disabled by default.

7. Set or override the following controller **DNS Server** configuration data:

- Name Servers** Provide a list of up to three DNS servers to forward DNS queries if the controller's DNS resources are unavailable. The DNS name servers are used to resolve IP addresses. Use the **Clear** link next to each DNS server to clear the DNS name server's IP address from the list.

Select **OK** to save the changes and overrides made to the DNS configuration. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's ARP Configuration

Overriding a Profile's Network Configuration

Address Resolution Protocol (ARP) is a protocol for mapping an IP address to a hardware MAC address recognized on the managed network. ARP provides protocol rules for making this correlation and providing address conversion in both directions. This ARP assignment can be overridden as needed, but removes the device configuration from the managed profile that may be shared with other similar device models.

When an incoming packet destined for a host arrives at the controller, the controller gateway uses ARP to find a physical host or MAC address that matches the IP address. ARP looks in its ARP cache and, if it finds the address, provides it so the packet can be converted to the right packet length and format and sent to the destination. If no entry is found for the IP address, ARP broadcasts a request packet in a special format to all the machines on the LAN to see if one machine knows that it has that IP address associated with it. A machine that recognizes the IP address as its own returns a reply so indicating. ARP updates the ARP cache for future reference and then sends the packet to the MAC address that replied.

To define an ARP supported configuration on the controller:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Network** to expand its sub menu options.
5. Select **ARP**.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

Switch VLAN Interface	IP Address	MAC Address	Device Type	
1	196.208.58.43	00-00-AE-33-00-00	Host	

Add Row

OK Reset Exit

FIGURE 131 Profile Overrides - Network ARP screen

6. Set or override the following parameters to define the controller's ARP configuration:

- Switch VLAN Interface** Use the spinner control to select a switch VLAN interface for an address requiring resolution.
- IP Address** Define the IP address used to fetch a MAC Address.
- MAC Address** Displays the target MAC address that's subject to resolution. This is the MAC used for mapping an IP address to a MAC address that's recognized on the managed network.
- Device Type** Specify the device type the ARP entry supports. Host is the default setting.

7. To add additional ARP overrides click on the **+ Add Row** button and enter the configuration information in the table above.
8. Select the **OK** button to save the changes and overrides to the ARP configuration. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's Quality of Service (QoS) Configuration

Overriding a Profile's Network Configuration

The controller uses different *Quality of Service (QoS)* screens to define WLAN and device radio QoS configurations for controller profiles.

QoS values are required to provide priority of service to some packets over others. For example, VoIP packets get higher priority than data packets to provide a better quality of service for high priority voice traffic.

The profile QoS screen maps the 6-bit *Differentiated Service Code Point* (DSCP) code points to the older

3-bit IP Precedent field located in the Type of Service byte of an IP header. DSCP is a protocol for specifying and controlling network traffic by class so that certain traffic types get precedence. DSCP specifies a specific per-hop behavior that is applied to a packet. This QoS assignment can be overridden as needed, but removes the device configuration from the managed profile that may be shared with other similar device models.

To define an QoS configuration for controller DSCP mappings:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Network** to expand its sub menu options.
5. Select **Quality of Service**.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

DSCP	802.1p Priority
0	0
1	0
2	0
3	0
4	0
5	0
6	0
7	0
8	1
9	1

FIGURE 132 Profile Overrides - Network QoS screen

6. Set or override the following parameters for IP DSCP mappings for untagged frames:

DSCP	Lists the DSCP value as a 6-bit parameter in the header of every IP packet used for packet classification.
802.1p Priority	Assign a 802.1p priority as a 3-bit IP precedence value in the Type of Service field of the IP header used to set the priority. The valid values for this field are 0-7. Up to 64 entries are permitted. The priority values are: 0 - <i>Best Effort</i> 1 - <i>Background</i> 2 - <i>Spare</i> 3 - <i>Excellent Effort</i> 4 - <i>Controlled Load</i> 5 - <i>Video</i> 6 - <i>Voice</i> 7 - <i>Network Control</i>

7. Use the spinner controls within the **802.1p Priority** field for each **DSCP** row to change or override the priority value.
8. Select the **OK** button located to save the changes and overrides. Select **Reset** to revert to the last saved configuration.

Overriding a Profile' Spanning Tree Configuration

Overriding a Profile's Network Configuration

Spanning Tree is a network layer protocol that ensures a loop-free topology in a mesh network of inter-connected layer 2 controllers. The spanning tree protocol disables redundant connections and uses the least costly path to maintain a connection between any two controllers in the network. Spanning tree protocol allows a network design that has one or more redundant links that provide a backup path if an active link fails. This switchover is automatic and does not require any human intervention.

Physical layer redundancy may also be provided using spanning tree. Spanning tree is a link management protocol that is part of the IEEE 802.1 standard for media access control bridges. Using the Dijkstra algorithm, STP provides link path redundancy between Ethernet devices while preventing undesirable loops in a network that can be created when multiple active paths exist between Ethernet controllers and bridges.

To establish path redundancy, STP creates a tree that spans all of the controllers in an extended network, forcing redundant paths into a blocked, state. STP allows only one active path at a time between any two network devices but establishes the redundant links as a backup if the preferred link should fail. If STP costs change, or if one network segment in the STP becomes unreachable, the spanning tree algorithm reconfigures the spanning tree topology and re-establishes the link by activating the standby path. Without spanning tree, multiple paths in the Ethernet network would be active resulting in an endless loop of traffic on the LAN.

Spanning Tree is supported on the RFS4000, RFS6000 and RFS7000 model controllers and can be used to provide link path redundancy when the controllers are connected to one or more external Ethernet switches. Spanning Tree can only support one active path per VLAN between Ethernet devices. If multiple paths per VLAN exist, redundant paths are blocked.

Multiple Spanning Tree Protocol (MSTP) is a VLAN-aware protocol and algorithm to create and maintain a loop-free network. It allows the configuration of multiple spanning tree instances. This ensures a loop-free topology for one or more VLANs. It allows the network administrator to provide a different path for each group of VLANs to better utilize redundancy.

Using MSTP, the network can be divided into regions. Each controller within a region uses the same VLAN to instance mapping. The entire network runs a spanning tree instance called the *Common Spanning Tree* instance (CST) that interconnects regions as well as legacy (STP and RSTP) bridges. The regions run on a local instance for each configured MSTP instance.

This Spanning Tree assignment can be overridden as needed, but removes the device configuration from the managed profile that may be shared with other similar device models.

To define or override a profile's spanning tree configuration:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.
2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.
3. Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.
4. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
5. Select **Network** to expand its sub menu options.
6. Select **Spanning Tree**.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

The screenshot displays the Network Spanning Tree configuration interface. It is divided into several sections:

- MSTP Configuration:**
 - MSTP Enable: (with a blue override icon)
 - Max Hop Count: 20 (range 7 to 127)
 - MST Config Name: My Name
 - MST Revision Level: 0 (range 0 to 255)
 - Cisco MSTP Interoperability: Disable
 - Hello Time: 2 (range 1 to 10)
 - Forward Delay: 15 (range 4 to 30)
 - Maximum Age: 20 (range 6 to 40)
- Spanning Tree Instance:**

Index	Priority	
6	32,768	<input type="checkbox"/>
<input type="button" value="Add Row"/>		
- PortFast:**
 - PortFast BPDU Filter: (with a blue override icon)
 - PortFast BPDU Guard: (with a blue override icon)
- Error Disable:**
 - Enable Recovery: (with a blue override icon)
 - Recovery Interval: 300 (range 10 to 1,000,000)
- VLANs:**

Index	VLANs	
1	4	<input type="checkbox"/>
<input type="button" value="Add Row"/>		

At the bottom of the screen are buttons for **OK**, **Reset**, and **Exit**.

FIGURE 133 Profile Overrides - Network Spanning Tree screen

7. Set or override the following **MSTP Configuration** parameters:

MSTP Enable	Enables the <i>Multiple Spanning Tree Protocol</i> (MSTP) feature. Select the check box to enable spanning tree for this device. This feature is disabled by default.
Max. Hop Count	Set the maximum number of hops used when creating a Spanning Tree. This value represents the maximum allowed hops for a BPDU (<i>Bridge Protocol Data Unit</i>) in an MSTP region. This value is used by all the MSTP instances. Enter a value between 7 - 127, or use the spinner control to set the value. The default setting is 20.
MST Config Name	Enter a name for the MST region. This is used when configuring multiple regions within the network. Each controller running MSTP is configured with a unique MST region name. This helps when keeping track of MSTP configuration changes. The name cannot exceed 64 characters.
MST Revision Level	Assign a MST revision level (0 - 255) to the MSTP region to which the device belongs. Each controller is configured with a unique MSTP name and revision number. This helps when keeping track of MSTP configuration changes. Increment this number with each configuration change. The revision level specifies the revision level of the current configuration. The default setting is 0.
Cisco MSTP Interoperability	Select <i>Enable</i> or <i>Disable</i> from the drop-down menu. This enables interoperability with Cisco's version of MSTP, which is incompatible with standard MSTP. The default setting is disabled.
Hello Time	The hello time is the time interval (in seconds) the device waits between BPDU transmissions. A low value leads to excessive traffic on the network, whereas a higher value delays the detection of a topology change. Set a hello time between 1 - 10 seconds. You can also use the spinner control next to the text-box to increase or decrease the value. The default setting is 2.
Forward Delay	The forward delay is the maximum time (in seconds) the root device waits before changing states (from a listening state to a learning state to a forwarding state). Set a value between 4 -30. You can also use the spinner control next to the text-box to increase or decrease the value. The default is 15.
Maximum Age	The max-age is the maximum time (in seconds) for which, if a bridge is the root bridge, a message is considered valid. This prevents frames from looping indefinitely. The max-age should be greater than twice the value of hello time plus one, but less than twice the value of forward delay minus one. Configure this value sufficiently high, so a frame generated by root can be propagated to the leaf nodes without exceeding the max age. Set the value between 6 - 40. You can also use the spinner control next to the text-box to increase or decrease the value. The default setting is 20.

8. Define or override the following **PortFast** configuration parameters:

PortFast BPDU Filter	Select the check box to enable BPDU filter for all portfast enabled ports. The Spanning Tree Protocol sends BPDUs from all the ports. Enabling the BPDU filter ensures PortFast enabled ports do not transmit or receive BPDUs.
PortFast BPDU Guard	Select the check box to enable BPDU guard for all portfast enabled ports. When the BPDU Guard feature is set for bridge, all portfast-enabled ports that have BPDU set to default shutdown the port upon receiving a BPDU. Thus no BPDUs are processed.

- Set or override the following **Error Disable** recovery parameters:

Enable Recovery	Select this check box to enable an error disable timeout caused by a BPDU guard. This option is disabled by default.
Recovery Interval	Define an interval (between 10 - 1,000,000) after which a recovering port is enabled. The default recovery interval is 300.

- Set or override the **Spanning Tree Instance** configuration.

Define a numerical index for each instance to assign each a unique priority. The **Priority** is assigned to an individual bridge based on whether it is selected as the root bridge. The lower the priority, the greater likelihood the bridge becoming the root for this instance.

- Use the **+ Add Row** button to create a new row in the table. To delete a row, select the row's delete icon.
- Refer to the **VLANs** table to associate a VLAN ID with the Instance index. You can add multiple VLANs to an instance.

Use the **+ Add Row** button to create a new row in the table. To delete a row, select the row's delete icon.
- Select **OK** to save or override the changes. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's Static Route Configuration

Overriding a Profile's Network Configuration

Use the **Static Routes** screen to set or override Destination IP and Gateway addresses enabling assignment of static IP addresses for requesting clients without creating numerous host pools with manual bindings. This eliminates the need for a long configuration file and reduces the controller resource space required to maintain address pools.

To create or override a profile's static routes:

- Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.
- Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.
- Select **Profile Overrides** from the Device menu to expand it into sub menu options.
- Select **Network** to expand its sub menu options.
- Select **Static Routes**.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

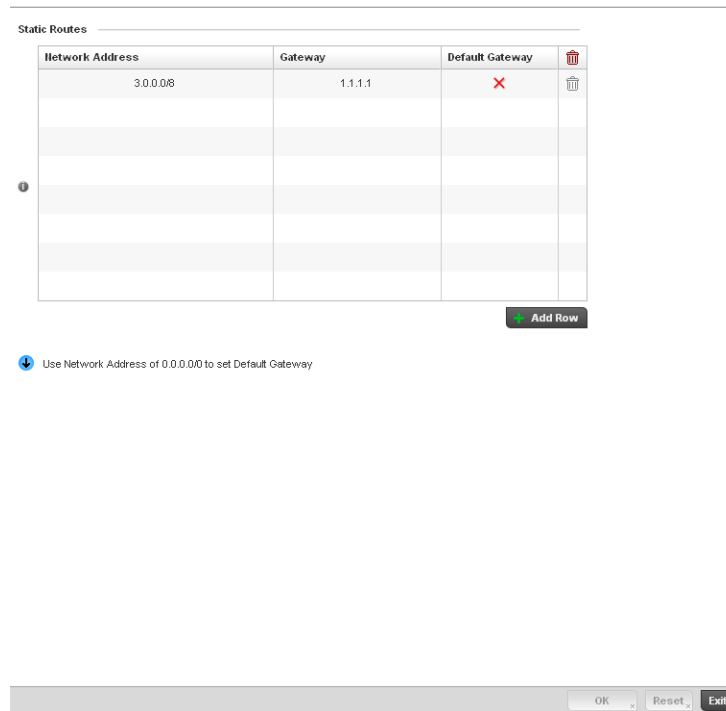


FIGURE 134 Static Routes screen

6. Select **Add Row +** as needed to include single rows in the static routes table.
7. Add IP addresses and network masks in the **Network** column.
8. Set or override the **Gateway** used to route traffic.

A green checkmark in the Default Gateway column defines a default gateway being applied. A red "X" means a gateway assignment has been made.

A default gateway is desirable when an IP address does not match any other routes in the controller's routing table. A gateway routes traffic from a managed device to another network segment. The controller's default gateway connects the managed network to the outside network (Internet). The gateway is associated with a router, which uses headers and forwarding tables to determine where packets are sent, and the controller, providing the path for the packet in and out of the gateway. Setting a default gateway for a device profile can help segregate network traffic, on behalf of a profile, to a single default gateway.

9. Select **OK** to save the changes and overrides. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's Forwarding Database Configuration

Overriding a Profile's Network Configuration

A *Forwarding Database* is used to forward or filter packets on behalf of the controller. The bridge reads the packet's destination MAC address and decides to either forward the packet or drop (filter) it. If it's determined the destination MAC is on a different network segment, it forwards the packet to the segment. If the destination MAC is on the same network segment, the packet is dropped (filtered). As nodes transmit packets through the bridge, the bridge updates its forwarding database with known MAC addresses and their locations on the network. This information is then used to decide to filter or forward the packet.

This forwarding database assignment can be overridden as needed, but removes the device configuration from the managed profile that may be shared with other similar device models.

To define or override a profile's forwarding database configuration:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Network** to expand its sub menu options.
5. Select **Forwarding Database**.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

The screenshot shows a configuration window titled "Aging Time". At the top, there is a "Bridge Aging Time" field with a value of 300 and a unit of "(0,10-1000000 seconds)". Below this is a "Static Forwarding Table" with three columns: "MAC Address", "VLAN Id", and "Interface Name". The table contains one row with the following data:

MAC Address	VLAN Id	Interface Name
01-02-03-04-AA-25	1	test

At the bottom right of the table is an "Add Row" button. At the bottom of the window are "OK", "Reset", and "Exit" buttons.

FIGURE 135 Profile Overrides - Network Forwarding Database screen

6. Define or override a **Bridge Aging Time** between 0, 10-1,000,000 seconds.

The aging time defines the length of time an entry will remain in the a bridge's forwarding table before being deleted due to lack of activity. If an entry replenishes a destination generating continuous traffic, this timeout value will never be invoked. However, if the destination becomes idle, the timeout value represents the length of time that must be exceeded before an entry is deleted from the forwarding table. The default setting is 300 seconds.
7. Use the **+ Add Row** button to create a new row within the MAC address table.
8. Set or override a destination **MAC Address**. The bridge reads the controller packet's destination MAC address and decides to forward the packet or drop (filter) it. If it's determined the destination MAC is on a different network, it forwards the packet to the segment. If the destination MAC is on the same network segment, the packet is dropped (filtered).
9. Define or override the target **VLAN ID** if the destination MAC is on a different network segment.
10. Provide an **Interface Name** used as the target destination interface for the target MAC address.
11. Select **OK** to save the changes and overrides. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's Bridge VLAN Configuration

Overriding a Profile's Network Configuration

A *Virtual LAN (VLAN)* is separately administrated virtual network within the same physical managed network. VLANs are broadcast domains defined within switches to allow control of broadcast, multicast, unicast, and unknown unicast within a Layer 2 device.

Administrators often need to route traffic to interoperate between different VLANs. Bridging VLANs are only for non-routable traffic, like tagged VLAN frames destined to some other device which will untag it. When a data frame is received on a port, the VLAN bridge determines the associated VLAN based on the port of reception. Using forwarding database information, the Bridge VLAN forwards the data frame on the appropriate port(s). VLAN's are useful to set separate networks to isolate some computers from others, without actually having to have separate cabling and Ethernet switches. Controllers can do this on their own, without need for the computer or other gear to know itself what VLAN it's on (this is called port-based VLAN, since it's assigned by port of the switch). Another common use is to put specialized devices like VoIP Phones on a separate network for easier configuration, administration, security, or quality of service.

Two main VLAN bridging modes are available in Brocade Mobility:

- **Tunnel Mode:** In tunnel mode, the traffic at the AP is always forwarded through the best path. The APs decides the best path to reach the destination and appropriately forwards the packets. Setting the VLAN to tunnel mode will ensure that packets are Bridge packets between local ethernet ports, any local radios, and tunnels to other APs and wireless controller.
- **Local Mode:** Local mode is typically configured in remote branch offices where traffic on remote private LAN segment needs to be bridged locally. Local mode implies that the wired and the wireless traffic are to be bridged locally.

To define a bridge VLAN configuration or override for a device profile:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Network** to expand its sub menu options.
5. Select **Bridge VLAN**.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

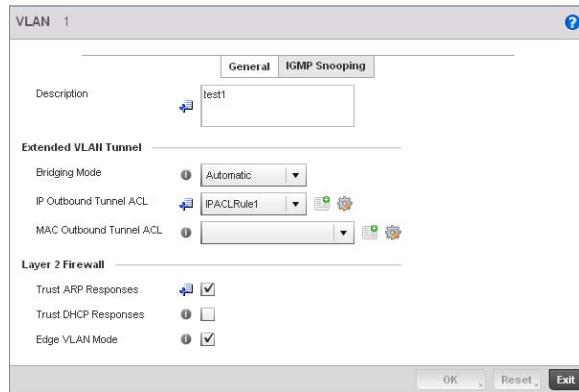


FIGURE 137 Profile Overrides - Network Bridge VLAN screen, General tab

8. The **General** tab displays by default.
9. If adding a new Bridge VLAN configuration, use the spinner control to define or override a **VLAN** ID between 1 - 4094. This value must be defined and saved before the General tab can become enabled and the remainder of the settings defined. VLAN IDs 0 and 4095 are reserved and unavailable.
10. Set or override the following General Bridge VLAN parameters:

Description	If creating a new Bridge VLAN, provide a description (up to 64 characters) unique to the VLAN's specific configuration to help differentiate it from other VLANs with similar configurations.
--------------------	---

11. Set or override the following **Extended VLAN Tunnel** parameters:

Bridging Mode	Specify one of the following bridging mode for use on the VLAN. <ul style="list-style-type: none"> • <i>Automatic</i>: Select automatic mode to let the controller determine the best bridging mode for the VLAN. • <i>Local</i>: Select Local to use local bridging mode for bridging traffic on the VLAN. • <i>Tunnel</i>: Select Tunnel to use a shared tunnel for bridging traffic on the VLAN. • <i>isolated-tunnel</i>: Select isolated-tunnel to use a dedicated tunnel for bridging traffic on the VLAN.
IP Outbound Tunnel ACL	Select an IP Outbound Tunnel ACL for outbound traffic from the pulldown menu. If an appropriate outbound IP ACL is not available click the create button to make a new one.
MAC Outbound Tunnel ACL	Select a MAC Outbound Tunnel ACL for outbound traffic from the pulldown menu. If an appropriate outbound MAC ACL is not available click the create button to make a new one.

NOTE

Local and Automatic bridging modes do not work with ACLs. ACLs can only be used with tunnel or isolated-tunnel modes.

12. Set or override the following **Layer 2 Firewall** parameters:

- Trust ARP Response** Select the checkbox to use trusted ARP packets to update the DHCP Snoop Table to prevent IP spoof and arp-cache poisoning attacks. This feature is disabled by default.
- Trust DHCP Responses** Select the checkbox to use DHCP packets from a DHCP server as trusted and permissible within the managed network. DHCP packets are used to update the DHCP Snoop Table to prevent IP spoof attacks. This feature is disabled by default.
- Edge VLAN Mode** Select the checkbox to enable edge VLAN mode. When selected, the edge controller's IP address in the VLAN is not used for normal operations, as its now designated to isolate devices and prevent connectivity. This feature is enabled by default.

13. Select the **OK** button to save the changes and overrides to the General tab. Select **Reset** to revert to the last saved configuration.

14. Select the **IGMP Snooping** tab to define or override the VLAN's IGMP configuration.

FIGURE 138 Profile Overrides - Network Bridge VLAN screen, IGMP Snooping

15. Define or override the following **IGMP Snooping** parameters for the Bridge VLAN configuration:

The *Internet Group Management Protocol* (IGMP) is a protocol used for managing members of IP multicast groups. The controller listens to IGMP network traffic and forwards the IGMP multicast packets to radios on which the interested hosts are connected. On the wired side of the network, the controller floods all the wired interfaces. This feature reduces unnecessary flooding of multicast traffic in the network.

- Enable IGMP Snooping** Select the check box to enable IGMP snooping on the controller. If disabled, snooping on a per VLAN basis is also disabled. This feature is enabled by default. If disabled, the settings under bridge configuration are overridden. For example, if IGMP snooping is disabled, but the bridge VLAN is enabled, the effective setting is disabled.
- Forward Unknown Multicast Packets** Select the check box to enable the controller to forward multicast packets from unregistered multicast groups. If disabled (the default setting), the unknown multicast forward feature is also disabled for individual VLANs.

16. Within the **Multicast Router** section, check the boxes of those interfaces used by the controller as a multicast router interface. Multiple controller interfaces can be selected and overridden.

Optionally select the **Snoop PIM-DVMRP Packets** box to snoop packets across the selected interface(s). This option is enabled by default.

17. Set or override the following **IGMP Querier** parameters for the profile's bridge VLAN configuration:

Enable IGMP Querier	Select the check box to enable IGMP querier. IGMP snoop querier is used to keep host memberships alive. It's primarily used in a network where there's a multicast streaming server and hosts subscribed to the server and no IGMP querier present. The controller can perform the IGMP querier role. An IGMP querier sends out periodic IGMP query packets. Interested hosts reply with an IGMP report packet. IGMP snooping is only conducted on wireless radios. IGMP multicast packets are flooded on wired ports. IGMP multicast packets are not flooded on the wired port. IGMP membership is also learnt on it and only if present, then it is forwarded on that port. A Brocade Mobility AP-7131 model access point can also be an IGMP querier.
Source IP Address	Define an IP address applied as the source address in the IGMP query packet. This address is used as the default VLAN querier IP address.
IGMP Version	Use the spinner control to set the IGMP version compatibility to either version 1, 2 or 3. The default setting is 3.
Maximum Response Time	Specify the maximum time (between 1 - 25 seconds) before sending a responding report. When no reports are received from a radio, radio information is removed from the snooping table. The controller only forwards multicast packets to radios present in the snooping table. For IGMP reports from wired ports, the controller forwards these reports to the multicast router ports. The default setting is 10 seconds.
Other Querier Timer Expiry	Specify an interval in either Seconds (60 - 300) or Minutes (1 - 5) used as a timeout interval for other querier resources. The default setting is 1 minute.

18. Select the **OK** button to save the changes and overrides to the IGMP Snooping tab. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's Cisco Discovery Protocol Configuration

Overriding a Profile's Network Configuration

The Cisco Discovery Protocol (CDP) is a proprietary Data Link Layer network protocol implemented in Cisco networking equipment and used to share information about network devices.

To override *Cisco Discovery Protocol* (CDP) configuration:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.
2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.
3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Network** to expand its sub menu options.

5. Select Cisco Discovery Protocol.

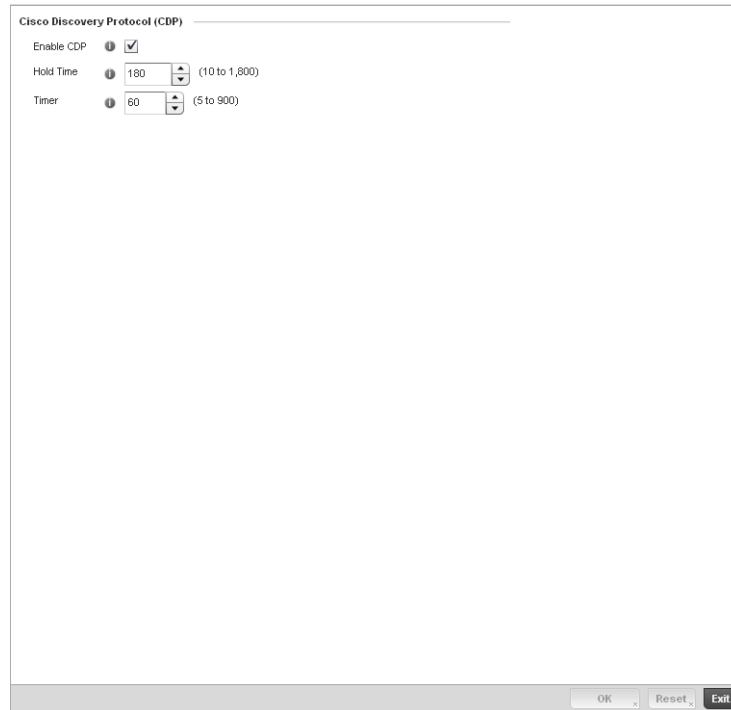


FIGURE 139 Profile Overrides - Network Cisco Discovery Protocol screen

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

6. Check the **Enable CDP** box to enable Cisco Discovery Protocol on the device.
7. Refer to the **Hold Time** field and use the spinner control to define a hold time between 10 - 1800 seconds for transmitted CDP Packets. The default value is 180 seconds.
8. Refer to the **Timer** field and use the spinner control to define a interval between 5 - 900 seconds to transmit CDP Packets. The default value is 60 seconds.
9. Select the **OK** button to save the changes and overrides. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's Link Layer Discovery Protocol Configuration

Overriding a Profile's Network Configuration

The Link Layer Discovery Protocol (LLDP) or IEEE 802.1AB is a vendor-neutral Data Link Layer protocol used by network devices for advertising of (announcing) their identity, capabilities, and interconnections

on a IEEE 802 LAN network. The protocol is formally referred to by the IEEE as Station and Media Access Control Connectivity Discovery. Both LLDP snooping and ability to generate and transmit LLDP packets will be provided.

Information obtained via CDP and LLDP snooping is available in the UI. In addition, information obtained via CDP / LLDP snooping is provided by an AP during the adoption process, so the L2 switch device name detected by the AP can be used as a criteria in the auto provisioning policy.

To override *Link Layer Discovery Protocol (LLDP)* configuration:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Network** to expand its sub menu options.
5. Select **Link Layer Discovery Protocol**.

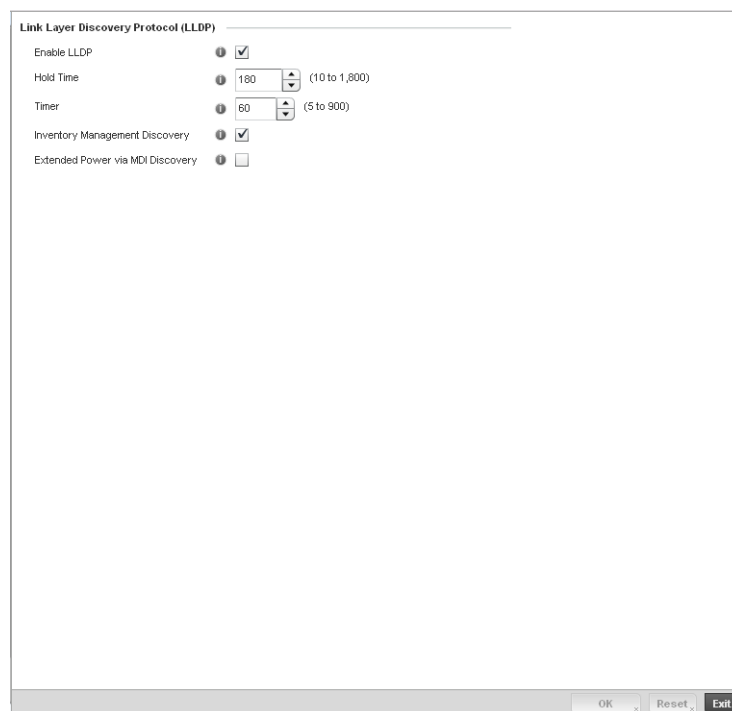


FIGURE 140 Profile Overrides - Network Link Layer Discovery Protocol screen

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

6. Check the **Enable LLDP** box to enable Link Layer Discovery Protocol on the device.
7. Refer to the **Hold Time** field and use the spinner control to define a hold time between 10 - 1800 seconds for transmitted LLDP Packets. The default value is 180 seconds.

8. Refer to the **Timer** field and use the spinner control to define the interval between 5 - 900 seconds to transmit LLDP Packets. The default value is 60 seconds.
9. Check the **Inventory Management Discovery** box to enable this feature. Inventory Management Discovery is used to track and identify inventory attributes including manufacturer, model, or software version.
10. Select the **Extended Power via MDI Discovery** box to enable this feature. Extended Power via MDI Discovery provides detailed power information through end points and other connected devices.
11. Select the **OK** button to save the changes and overrides. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's Miscellaneous Network Configuration

Overriding a Profile's Network Configuration

A controller profile can be configured to include a hostname in a DHCP lease for a requesting device and its profile. This helps an administrator track the leased DHCP IP address by hostname for the controller supported device profile. When numerous DHCP leases are assigned, an administrator can better track the leases when hostnames are used instead of devices.

To include a hostnames in DHCP request:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.
2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.
3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Network** to expand its sub menu options.
5. Select **Miscellaneous**.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

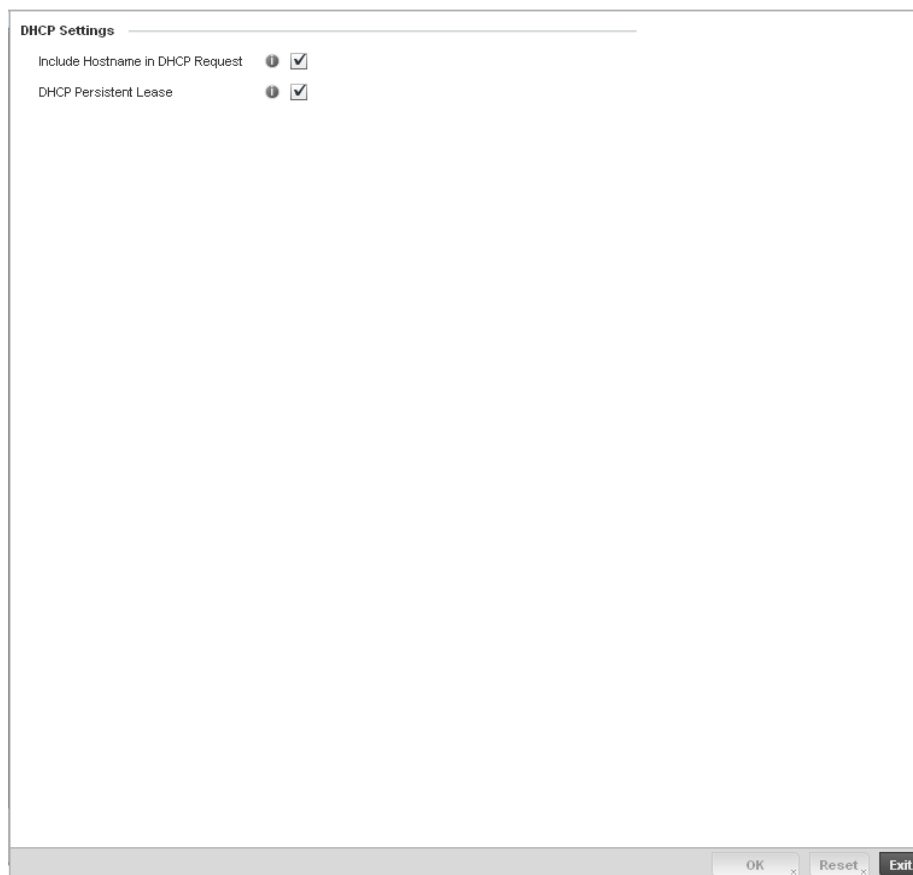


FIGURE 141 Profile Overrides - Network Miscellaneous screen

6. Refer to the DHCP Settings section to configure miscellaneous DHCP Settings.

Include Hostname in DHCP Request

Select the Include Hostname in DHCP Request checkbox to include a hostname in a DHCP lease for a requesting device. This feature is disabled by default.

DHCP Persistent Lease

Check this box to enable a persistent DHCP lease for the device. A persistent DHCP lease assigns the same IP Address and other network information to the device each time it renews its DHCP lease.

7. To enable critical resource monitoring for the device, select a **Critical Resource Policy** from the pull-down menu in the **Critical Resource Monitoring** section. If a new critical resource monitoring policy is needed click the **Create** button and specify the Ping Interval, IP Address, Ping Mode and VLAN for the devices being monitored.
8. Select the **OK** button to save the changes and overrides. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's Security Configuration

A controller or Access Point profile can have its own firewall policy, wireless client role policy, WEP shared key authentication, NAT policy and VPN policy (controller only) applied. If an existing firewall, client role or NAT policy is unavailable, an administrator can be navigated from the **Configuration > Profiles** section of the controller UI to the **Configuration > Security** portion of the UI to create the

required security policy configuration. Once created, a policy's configuration can have an override applied to meet the changing data protection requirements of a device's environment. However, in doing so the device must now be managed separately from the profile configuration shared by other devices within the managed network.

For more information on applying an override to an existing device profile, refer to the following sections:

- [Overriding a Profile's General Security Settings](#)
- [Overriding a Profile's Certificate Revocation List \(CRL\) Configuration](#)
- [Overriding a Profile's ISAKMP Configuration](#)
- [Overriding a Profile's Transform Set Configuration](#)
- [Overriding a Profile's VPN Configuration](#)
- [Overriding a Profile's NAT Configuration](#)

Overriding a Profile's General Security Settings

Overriding a Profile's Security Configuration

A controller profile can leverage existing firewall, wireless client role and WIPS policies and apply them to the profile's configuration. This affords each controller profile a truly unique combination of data protection policies best meeting the data protection requirements of that controller profile. However, as deployment requirements arise, an individual device may need some or all of its general security configuration overridden from that defined in the profile.

To configure a profile's security settings and overrides:

1. Select **Devices** from the Configuration tab.
The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.
2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.
Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.
3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Security** to expand its sub menu options.
5. Select **General**.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

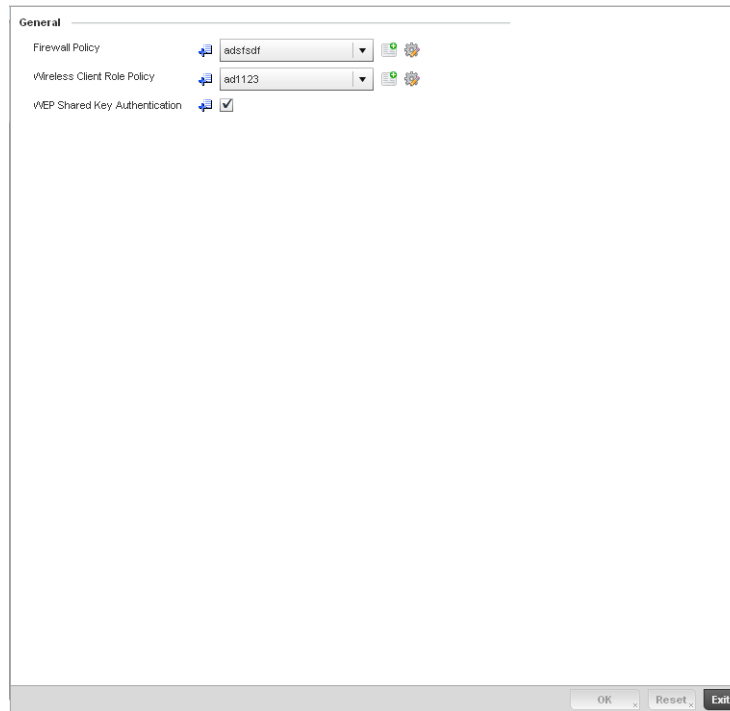


FIGURE 142 Profile Overrides - General Security screen

6. Refer to the **General** field to assign or override the following:

Firewall Policy

Use the drop-down menu to select an existing Firewall Policy to use as an additional security mechanism with this controller profile. All devices using this controller profile and Access Point must meet the requirements of the firewall policy to access the managed network. A firewall is a mechanism enforcing access control, and is considered a first line of defense in protecting proprietary information within the wireless controller managed network. The means by which this is accomplished varies, but in principle, a firewall can be thought of as mechanisms both blocking and permitting data traffic within the wireless controller managed network. If an existing Firewall policy does not meet your requirements, select the **Create** icon to create a new firewall policy that can be applied to this profile. An existing policy can also be selected and overridden as needed using the **Edit** icon. For more information, see *Wireless Firewall* on page 9-419 and *Configuring a Firewall Policy* on page 9-420.

Wireless Client Role Policy

Use the drop-down menu to select a client role policy the controller uses to strategically filter client connections based on a pre-defined set of filter rules and connection criteria. If an existing Wireless Client Role policy does not meet your requirements, select the **Create** icon to create a new configuration that can be applied to this profile. An existing policy can also be selected and overridden as needed using the **Edit** icon. For more information, see *Wireless Client Roles* on page 9-438.

WEP Shared Key Authentication

Select the check box to require devices using this profile to use a WEP key to access the managed network using this profile. The wireless controller, other proprietary routers, and clients use the key algorithm to convert an ASCII string to the same hexadecimal number. Clients without adapters need to use WEP keys manually configured as hexadecimal numbers. This option is disabled by default.

7. Select an **Advanced WIPS Policy** from the drop-down menu in the **Wireless IDS/IPS** section. Define an advanced WIPS configuration to optionally remove (terminate) unwanted device connections, and sanction (allow) or unsanction (disallow) specific events within the managed network.

If an existing Advanced WIPS policy does not meet the profile's data protection requirements, select the **Create** icon to create a new configuration that can be applied to the profile. An existing policy can also be selected and overridden as needed using the **Edit** icon. For more information, see [Configuring an Advanced WIPS Policy on page 9-456](#).

8. Select **OK** to save the changes or overrides. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's Certificate Revocation List (CRL) Configuration

Overriding a Profile's Security Configuration

A *certificate revocation list* (CRL) is a list of certificates that have been revoked or are no longer valid. A certificate can be revoked if the *certificate authority* (CA) had improperly issued a certificate, or if a private-key is compromised. The most common reason for revocation is the user no longer being in sole possession of the private key.

To define a Certificate Revocation configuration or override:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Security** to expand its sub menu options.
5. Select **Certificate Revocation**.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

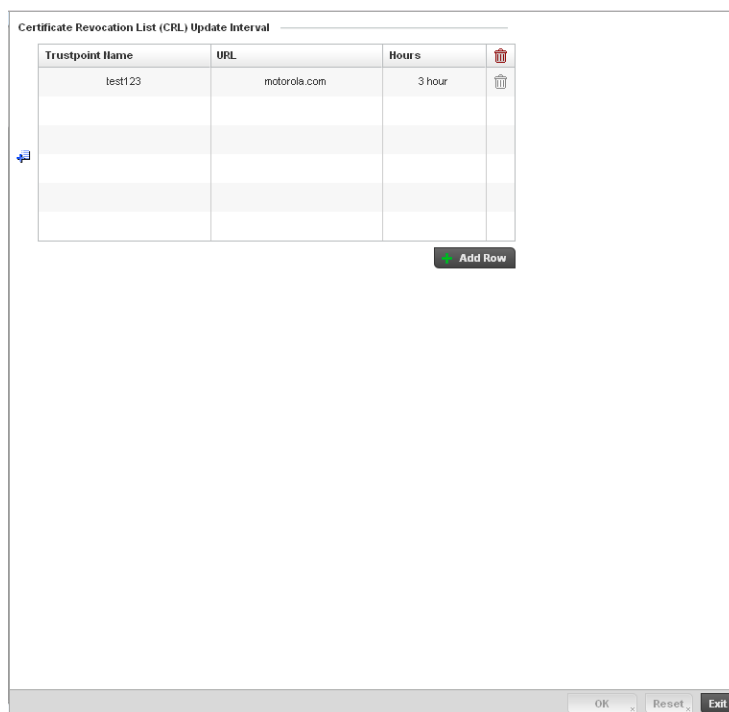


FIGURE 143 Profile Overrides - Certificate Revocation screen

6. Select the **+ Add Row** button to add a column within the **Certificate Revocation List (CRL) Update Interval** table to quarantine certificates from use in the managed network.

Additionally, a certificate can be placed on hold for a user defined period. If, for instance, a private key was found and nobody had access to it, its status could be reinstated.

 - a. Provide the name of the trustpoint in question within the **Trustpoint Name** field. The name cannot exceed 32 characters.
 - b. Enter the resource ensuring the trustpoint's legitimacy within the **URL** field.
 - c. Use the spinner control to specify an interval (in hours) after which a device copies a CRL file from an external server and associates it with a trustpoint.
7. Select **OK** to save the changes and overrides made within the Certificate Revocation screen. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's ISAKMP Configuration

Overriding a Profile's Security Configuration

ISAKMP (also known as IKE) is the negotiation protocol enabling two hosts to agree on how to build an IPSec security association. To configure the security appliance for virtual private networks, set global parameters that apply system wide and define policies peers negotiate to establish a VPN tunnel.

The ISAKMP protocol is an IPSec standard protocol used to ensure security for VPN negotiation, and remote host or network access. ISAKMP provides an automatic means of negotiation and authentication for communication between two or more parties. ISAKMP manages IPSec keys automatically.

The ISAKMP screen displays by default. The ISAKMP screen lists those ISAKMP policies created thus far. Use the ISAKMP Policy screen to configure the Internet Security Association and Key Management Protocol (ISAKMP) for creating a VPN. ISAKMP is a framework for authentication and key exchange. It defines the procedures and packet formats to establish, negotiate, modify and delete Security Associations (SA). Any of these policies can be selected and applied to the controller.

A VPN tunnel is negotiated in two phases. The first phase creates an ISAKMP SA that's a control channel. The data channel is negotiated using this control channel. ISAKMP policy parameters are not negotiated and the transform set is for negotiating the data channel (IPsec SAs).

To define an ISAKMP configuration or apply overrides:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Security** to expand its sub menu options.
5. Select **ISAKMP Policy**.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

ISAKMP Policy	Encryption	Hash Algorithm	Authentication Type	SA Lifetime	DH Group Identifier
abcf34	aes-256	SHA	Pre-Shared	1d 0h 0m 0s	2

FIGURE 144 Profile Overrides - ISAKMP Policy screen

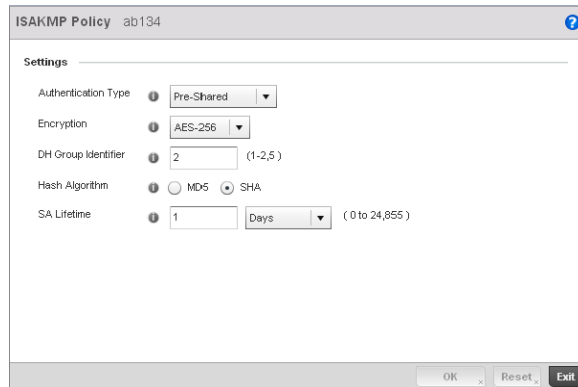
6. Refer to the following configuration data for existing ISAKMP policies:

ISAKMP Policy	Displays the name assigned to the ISAKMP policy when it was initially created. The name cannot be modified as part of the edit process.
Encryption	Lists the encryption type used for encrypting packets with this ISAKMP policy.
Hash Algorithm	Displays the MD5 or SHA hashing algorithm used in the ISAKMP policy.
Authentication Type	Lists the key sharing mechanism used for establishing a secure connection between two peers using this ISAKMP policy.
SA Lifetime	Lists the lifetime in either seconds, minutes, hours or days for the <i>security association</i> (SA) used by this ISAKMP policy.
DH Group Identifier	Displays the <i>Diffie-Hellman</i> (DH) group identifier used by this ISAKMP policy. DH is a cryptographic protocol that allows 2 entities that have no prior knowledge of each other to jointly derive and establish a shared secret key over an unsecured communication channel. This secret key can then be used to initiate a secure connection between the two entities.

7. Select **Add** to create a new ISAKMP policy, **Edit** to modify or override the attributes of a selected existing policy or **Delete** to remove obsolete policies from the list of those available to the controller.

If adding or editing an existing ISAKMP policy, the following screen displays.

When ISAKMP negotiations begin, the peer initiating the negotiation sends its policies to the remote peer. The remote peer searches for a match with its own policies using the defined priority scheme. A ISAKMP policy matches when they have the same encryption, hash, authentication and Diffie-Hellman settings. The ISAKMP lifetime must also be less than or equal to the lifetime in the policy sent. If the lifetimes do not match, the shorter lifetime applies. If no match exists, ISAKMP refuses negotiation.



The screenshot shows a window titled "ISAKMP Policy ab134" with a "Settings" section. The settings are as follows:

Setting	Value	Range/Options
Authentication Type	Pre-Shared	Dropdown
Encryption	AES-256	Dropdown
DH Group Identifier	2	(1-2,5)
Hash Algorithm	SHA	Radio buttons: MD5, SHA
SA Lifetime	1	Days (0 to 24,855)

At the bottom of the window are buttons for "OK", "Reset", and "Exit".

FIGURE 145 ISAKMP Policy Add screen

8. Set or override the following configuration parameters for the new or modified ISAKMP policy:

ISAKMP Policy	If creating a new ISAKMP policy, assign it a name to help differentiate it from others that may have a similar configuration. The policy name cannot exceed 64 characters. The name cannot be modified as part of the edit process.
Authentication Type	<p>Set the key sharing mechanism used for establishing a secure connection between two peers using this ISAKMP policy. Use the drop-down menu to select one of the following two authentication options:</p> <p><i>Pre-Shared</i> – Select this option to use a key shared between the VPN endpoints. Pre-Shared is the default setting.</p> <p><i>RSA Signature</i> – Uses a RSA signature as the authentication key. Ensure digital certificates and RSA keys have been installed on the target system before using this option.</p>
Encryption	<p>Use the drop-down menu to select the encryption algorithm to use. Select from the drop-down list.</p> <p><i>DES</i> – DES stands for <i>Data Encryption Standard</i>. It uses a 56-bit key for encryption. This standard is deprecated and replaced by the 3DES standard. It is provided for backward compatibility.</p> <p><i>3DES</i> - 3DES or <i>Triple DES</i> is an encryption standard that replaced DES. It provides a simple method of increasing the key size of the DES algorithm to protect from brute force attacks. It uses a set of three (3) standard 56-bit DES keys to provide increased key length for encryption.</p> <p><i>AES</i> – AES stands for <i>Advanced Encryption Standard</i> or the Rijndael Encryption Algorithm that was adopted as the new FIPS standard in the year 2002. It is a symmetric-key encryption standard that uses three (3) block ciphers of length 128, 192, 256 bits. This option represents the AES-128 bit block cipher.</p> <p><i>AES-192</i> – AES stands for <i>Advanced Encryption Standard</i> or the Rijndael Encryption Algorithm that was adopted as the new FIPS standard in the year 2002. It is a symmetric-key encryption standard that uses three (3) block ciphers of length 128, 192, 256 bits. This option represents the AES-192 bit block cipher.</p> <p><i>AES-256</i> – AES stands for <i>Advanced Encryption Standard</i> or the Rijndael Encryption Algorithm that was adopted as the new FIPS standard in the year 2002. It is a symmetric-key encryption standard that uses three (3) block ciphers of length 128, 192, 256 bits. This option represents the AES-256 bit block cipher. AES -256 is the default setting.</p>
DH Group Identifier	<p>Set the <i>Diffie-Hellman</i> (DH) group identifier used by this ISAKMP policy. DH is a cryptographic protocol that allows 2 entities that have no prior knowledge of each other to jointly derive and establish a shared secret key over an unsecured communication channel. This secret key can then be used to initiate a secure connection between the two entities.</p> <p>The valid values are 1, 2 and 5 and indicates the group used for the key exchange. The default setting is 2.</p>
Hash Value	<p>Set the hash algorithm. Select from:</p> <p><i>MD5</i> – MD5 or <i>Message-Digest algorithm 5</i> is a popular 128-bit hash-function. It is commonly used for checking the integrity of files.</p> <p><i>SHA</i> – <i>Secure Hash Algorithm</i> (SHA) is a NIST certified FIPS hash algorithm. SHA is the default setting.</p>
SA Lifetime	<p>Set the lifetime in seconds for the <i>security association</i> (SA) used by this ISAKMP policy. This is the lifetime of the ISAKMP SA. The lifetime for ESP/AH SAs are configured separately.</p> <p><i>Days</i> – Sets the SA duration in days (1 - 24,856).</p> <p><i>Hours</i> – Sets the SA duration in hours (1 - 596,524).</p> <p><i>Minutes</i> – Sets the SA duration in minutes (1 - 35, 791, 395).</p> <p><i>Seconds</i> – Sets the SA duration in seconds (60 - 2,147,483,646). The default setting is 86,400 seconds.</p>

9. Select **OK** to save the changes or overrides. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's Transform Set Configuration

Overriding a Profile's Security Configuration

Use the **Transform Set** screen to configure and manage transform sets. A Transform Set is a set of parameters that transform an IP packet from clear text to cipher text. The transform set is an acceptable combination of security protocols, algorithms and other settings that are applied to IPSec protected traffic.

With manually established security associations, there's no negotiation with the peer. Both sides must specify the same transform set, regardless of whether the SA is manual or automatic. For manual SAs, the ISAKMP policy does not apply. If you change a Transform Set definition, the change is only applied to Crypto Map entries that reference the Transform Set. If a transform-set is changed, the existing SAs are removed.

To define a transform set configuration or override that can be applied to a controller profile:

1. Select **Devices** from the Configuration tab.
The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.
2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.
Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.
3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Security** to expand its sub menu options.
5. Select **Transform Set**.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

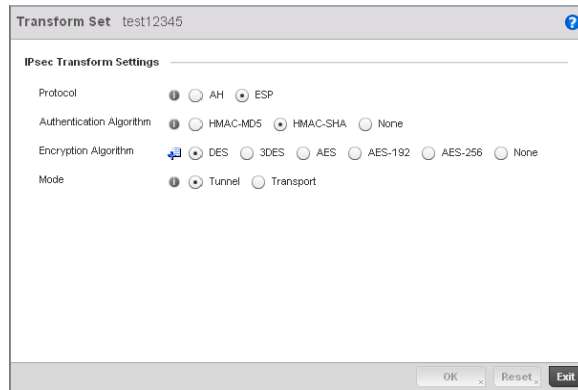


FIGURE 147 Transform Set Configuration screen

8. Set or override the following configuration parameters for the Transform Set:

Transform Set	If creating a new Transform Set, assign it a name to help differentiate it from others that may have a similar configuration. The name cannot exceed 64 characters. The name cannot be modified as part of the edit process.
Protocol	Select the radio button of the IPSec protocol used with the Transform Set. AH provides data authentication only. ESP provides data confidentiality and well as authentication. ESP is the default setting.
Authentication Algorithm	Set the authentication algorithm used to validate identity. <i>HMAC-MD5</i> – Use the <i>Message Digest 5 (MD5)</i> as the HMAC algorithm. <i>HMAC-SHA</i> – Use the <i>Secure Hash Algorithm (SHA)</i> as the HMAC algorithm. HMAC-SHA is the default setting. <i>None</i> – Applies no authentication. If the protocol is AH, None cannot be selected.
Encryption Algorithm	The encryption algorithm radio button only displays when ESP is selected as the Transform Set protocol. By default, the Transform set uses AES-256. It's a symmetric-key encryption standard that uses a block ciphers length 256 bits. AES -256 is the default setting. Selecting <i>None</i> applies no encryption. When the protocol is ESP, encryption and authentication cannot both be set to None.
Mode	Set the mode used for packet organization with respect to header location and the scope of ESP or AH protection boundary. Use Tunnel for site-to-site VPN and Transport mode for remote VPN configurations. The default mode is Tunnel.

9. Select **OK** to save the updated or overrides. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's VPN Configuration

Overriding a Profile's Security Configuration

A *Virtual Private Network (VPN)* provides secure communications between two or more IP networks or computer devices across an IP network infrastructure. VPN is commonly deployed to securely inter-connect remote sites or provide remote access across the public Internet and connect networks across a private IP network. Traffic exchanged between VPN end-points can transparently pass through shared resources such as routers, and other telecommunications equipment and is secured using *IP Security (IPSec)*.

To define a VPN configuration or override that can be applied to a controller profile:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Security** to expand its sub menu options.
5. Select **VPN Configuration**.

The screenshot shows the 'Basic Settings' tab for VPN Configuration. It features two main sections: 'Peer Settings' and 'Aggressive Mode Peer (ISAKMP Responder)'. The 'Peer Settings' section contains a table with columns for 'Peer Address' and 'Key'. The first row shows '192.167.81.57' and '*****'. Below this table is an 'Add Row' button. The 'Additional Settings' section includes fields for 'ISAKMP Keepalive' (set to 10), 'IPsec Lifetime (seconds)' (set to 1, with a unit dropdown set to 'Hours'), and 'IPsec Lifetime (kB)' (set to 4608000). The 'Aggressive Mode Peer (ISAKMP Responder)' section contains a table with columns for 'Peer Type', 'Peer Id', and 'Preshared Key'. The first row shows 'IP Address', '192.169.51.45', and '*****'. Below this table is another 'Add Row' button. At the bottom of the window are 'OK', 'Reset', and 'Exit' buttons.

FIGURE 148 Profile Overrides - VPN Configuration, Basic Settings tab

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

6. The **Basic Settings** tab displays by default. Refer to the Peer Settings table to add peer addresses and keys for VPN tunnel destinations. Use the **+ Add Row** function as needed to add additional destinations and keys.
7. Set or override the **Peer Address** from (within the **Peer Settings** table) for a known device at the other end of the managed VPN tunnel.
8. Define or override the **Key** (between 8 - 21 characters long) shared between the peer and a device (and is required by devices at both ends of the tunnel for inter operation). The key is exposed as a series of "*****" characters unless the **Show** radio button is selected to display the actual characters comprising the key. This is the key set for policy.

9. Select the **More >** button to the right of the **Additional Settings** field to expand the screen to display several IPsec parameters and the **Aggressive Mode Peer** table.

10. Set or override the following IPsec parameters within the **Additional Settings** field:

IPsec ISAKMP Keep Alive	Use the spinner control to set the IPsec ISAKMP keep alive duration (between 10 - 3,600) seconds. This determines the frequency of IKE keep alive messages for dead peer detection. The default value is 10 seconds.
IPsec Lifetime (seconds)	Set the IPsec Lifetime value in either Seconds (90 - 2,147,483,646), Minutes (2 - 35,791,395), Hours (1 - 596,524) or Days (1 - 24,856). This value represents the time-based IPsec security association lifetime. The default setting is 3,600 seconds. The IPsec Lifetime this is a global setting applying to all VPN Security associations. This global lifetime value can be overridden by crypto maps. The default value is 10 seconds.
IPsec Lifetime (kB)	Sets the IPsec Lifetime value in kilo bytes (500 - 2,147,483, 646). The life time of the VPN tunnel is decided by the amount of traffic that passes through the VPN. Once this value is exceeded, the VPN is automatically closed and the clients have to re-negotiate the security association. The default is 4,608,000 kilobytes. The IPsec Lifetime this is a global setting applying to all VPN Security associations. This global lifetime value can be overridden by crypto maps.

NOTE

The security association is re-negotiated when the IPsec Lifetime value expires, whether it's seconds or kB, whichever comes first.

11. Refer to the **Aggressive Mode Peer (ISAKMP Responder)** table to set or override the following parameters:

Aggressive mode is an IKE policy mode which reduces the number of steps during IKE authentication negotiation between two IKE peers. This mode has less negotiation power and does not provide identity protection between the peers. This mode is vulnerable to man-in-the-middle attacks.

12. Select the **+ Add Row** button to add additional rows.

Peer Type	Define the type of aggressive mode peer used as one of the following. <i>IP Address</i> - The IP address of the peer is provided. <i>Hostname</i> - The Host Name of the peer is provided. <i>Distinguished Name</i> - The peer type is defined by a distinguished name.
Peer ID	Set the Peer Id as either an IP address, hostname or distinguished name.
Preshared Key	Provide a shared key between this device and the aggressive mode peer. Both use the same key to interoperate. The key is exposed as a series of "*****" characters unless the Show radio button is selected to display the actual characters comprising the key.

13. Select **OK** to save the changes or overrides. Select **Reset** to revert to the last saved configuration.

14. Select the **Crypto Map** tab.

Crypto Map entries are sets of configuration parameters for encrypting packets that pass through the VPN Tunnel. The screen lists the various Crypto Map entries defined under a particular VPN Policy.

15. Either select **Add** to create a new Crypto Map configuration, **Edit** to modify or override the attributes of an existing Crypto Map configuration or **Delete** to permanently remove a selected configuration.
16. Provide a **Name** for the Crypto Map up to 64 characters in length. The name should help distinguish the Crypto Map from others with similar configurations. Select **Continue**.
17. Select **Add** to display a **Crypto Map Entry** screen where the attributes of the Crypto Map configuration can be defined or overridden.

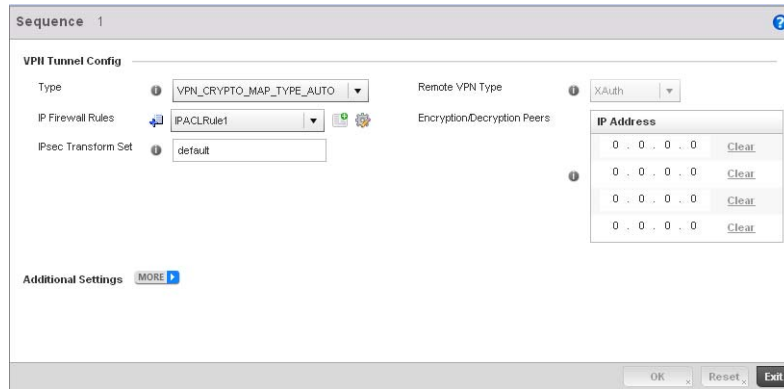


FIGURE 149 Crypto Map Entry screen

18. Use the spinner control to set or override a **Sequence** between 1 - 1000 that identifies the Crypto Map in the list of Crypto Map entries.
19. Define or override the following **VPN Tunnel Config** parameters:

Type Define the type of tunnel to be created. Select from one of the following:
site-to-site-auto – Sets a VPN connection (using ISAKMP) to a remote site where the connection parameters are configured automatically. This is the default setting.
remote-vpn – Creates a VPN connection to a remote client.
site-to-site-manual – Creates a VPN connection to a remote site where the connection parameters are configured manually.

IP Firewall Rules Set the Firewall rules applied to this VPN tunnel. Select existing rules from the drop-down list. These rules are used to classify traffic for VPN.
 Use the **Create** icon to create a new Firewall rule set or the **Edit** icon to modify the configuration of an existing Firewall rule. For more information, see [Configuring IP Firewall Rules on page 9-431](#).

IPsec Transform Set Define the existing (previously configured) transform set for this VPN Tunnel.

Remote VPN Client Type This parameter is available only when *Remote VPN* is selected in the **Type** field. Select from one of the following:
IPsec-L2TP – The remote VPN client uses an IPsec Layer 2 Tunnelling Protocol.
XAuth – The remote VPN client uses X-Windows Authorization.

Encryption/Decryption Peers If *Automatic Site-to-Site VPN* or *Remote VPN Client* is selected as the Type, Enter up to 4 IP addresses for Encryption/Decryption Peers. For Auto VPN, multiple peers are configured for redundancy. An administrator is required to configure just one peer, unless redundancy is required and another peer is needed.

20. Select the **More >** link to expand the screen to add the following **Additional Settings** or overrides for the Crypto Map configuration:

Perfect Forward Secrecy (PFS)	Select the check box to enable PFS and set DH group or either 1, 2 or 5. Enable this option to set an authenticated key-agreement protocol that uses public key cryptography. <i>Perfect Forward Secrecy</i> (PFS) ensures a session key derived from public and private keys is not compromised if one of the private keys is compromised in the future. This feature is disabled by default.
Mode of Tunnel Initiation	Select the tunnel mode. Choose from: <i>Main</i> – The complete process of authentication and key exchange is followed when the Main mode is selected. Main is the default value. <i>Aggressive</i> – A truncated process of authentication and key exchange is followed when Aggressive mode is selected.
Local Identity Type (ISAKMP Initiator)	Define how this end of the tunnel is identified. Select from: <i>IP Address</i> – This end of the tunnel is identified by an IP address. IP Address is the default setting. <i>Hostname (FQDN)</i> – This end of the tunnel is identified by a host name. <i>Distinguished Name</i> – This end of the tunnel is identified by a distinguished name.
Local Identity Value	Set the name of the host or domain (only applicable for the initiator). The name cannot exceed 64 characters. This option is not available when IP Address is the Local Identity Type.

21. Set or override the following **IPsec Security Association (SA)** parameters for the Crypto Map configuration:

Lifetime (seconds)	Select the check box to enable the spinner control to set the Security Association lifetime in seconds (90 - 2,147,483,646). This feature is disabled by default. This is the override for the value defined within the Basic Settings tab.
Lifetime (kb)	Select the check box to enable the spinner control to set the Security Association lifetime in kilobytes (500 - 2,147,483,646) for data traversing the VPN. This feature is disabled by default. This is the override for the value defined within the Basic Settings tab.
Per Host SA	Select the check box to set the granularity for IPsec security associations to one per-host. This feature is disabled by default.

22. Select **OK** to save the updates and overrides. Select **Reset** to revert to the last saved configuration.

23. Select the **Remote VPN** tab to define an additional remote VPN configuration beyond what has been defined in the Basic Settings tab.

The screenshot shows the 'Remote VPN' configuration window with three tabs: 'Basic Settings', 'Crypto Map', and 'Remote VPN'. The 'Remote VPN' tab is active.

Local Pool Settings

Low Address	High Address	
192.168.1.81	192.168.1.128	

Remote VPN Config

DNS Server: 192.168.1.1
WINS Server: 192.168.1.2

Remote VPN Authentication

Authentication Method: RADIUS Local None

Local Users Config

Name	Password	
admin	*****	

Buttons: OK, Reset, Exit

FIGURE 150 Remote VPN screen

24. Use the **Local Pool Settings** table set or override a **Low Address** and **High Address** range for local IP address assignments from the local pool. This table represents the virtual IP pool that can be assigned to remote VPN clients. Select the **+ Add Row** button as required to additional entries in the table.
25. Refer to the **Remote VPN Config** field to set or override the following parameters:
 - DNS Server** Enter the address of the DNS Server configured on the remote VPN client.
 - WINS Server** Enter the address of the WINS Server configured on the remote VPN client.
26. Select the radio button corresponding to the desired **Authentication Method** used to authenticate the remote VPN client. Set to **RADIUS** to use a dedicated RADIUS Server (If using XAuth as the remote VPN Client Type) or **Local** for internal controller resources for authentication. Selecting **None** bypasses authenticating the remote client. Local is the default setting.
27. If using the Local default setting, select **+ Add Row** to add a user name for the local login account and the password for the user name on the local server. The password is exposed as a series of "*****" characters unless the **Show** check box is selected to display the actual characters comprising the password.
28. If selecting RADIUS, specify the **NAS** originating the RADIUS accept request and define a RADIUS configuration for remote client authentication within the RADIUS Configuration field.
29. If selecting RADIUS, specify an **ARP Timeout** value in *Days, Hours, Minutes* or *Seconds*.

RADIUS Servers

Server Type	Server IP	Authentication Port	Password	
★ Primary ★	234. 54 . 9 . 55	<input checked="" type="checkbox"/> 1024	★ ***** <input type="checkbox"/> Show	

Add Row

FIGURE 151 Remote VPN Authentication screen - RADIUS Configuration

30. Set or override the following **RADIUS Configuration** parameters:

- Server Type** Select the RADIUS server type. Select from:
Primary – The selected server is the Primary RADIUS Server.
Secondary – The selected server is the Secondary RADIUS Server.
- Server IP** Set IP address of the RADIUS server.
- Authentication Port** Select the check box and use the spinner control to set the UDP port number where the defines RADIUS server is listening.
- Password** Enter the password to log on to the RADIUS server. The password is exposed as a series of “*****” characters unless the **Show** check box is selected to display the actual characters comprising the password.

31. Select **OK** to update the Remote VPN Authentication settings and overrides. Select **Reset** to revert to the last saved configuration.

Overriding a Profile’s NAT Configuration

Overriding a Profile’s Security Configuration

Network Address Translation (NAT) is a technique to modify network address information within IP packet headers in transit across a traffic routing device. This enables mapping one IP address to another to protect wireless controller managed network address credentials. With typical deployments, NAT is used as an IP masquerading technique to hide private IP addresses behind a single, public facing, IP address.

NAT is a process of modifying network address information in IP packet headers while in transit across a traffic routing device for the purpose of remapping one IP address to another. In most deployments NAT is used in conjunction with IP masquerading which hides RFC1918 private IP addresses behind a single public IP address.

NAT can provide a controller profile outbound Internet access to wired and wireless hosts connected to either a thick access point (such as an AP71xx, AP71xx or AP-5131 model) or a RFS4000, RFS6000 or RFS7000 model wireless controller. Many-to-one NAT is the most common NAT technique for outbound Internet access. Many-to-one NAT allows a thick access point or wireless controller to translate one or more internal private IP addresses to a single, public facing, IP address assigned to a 10/100/1000 Ethernet port or 3G card.

To define a NAT configuration or override that can be applied to a controller profile:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Security** to expand its sub menu options.
5. Select **NAT**.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.




FIGURE 152 Profile Overrides - NAT Pool screen

6. The **NAT Pool** displays by default. The NAT Pool screen lists those NAT policies created thus far. Any of these policies can be selected and applied to a controller profile.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

7. Select **Add** to create a new NAT policy that can be applied to a controller profile. Select **Edit** to modify or override the attributes of an existing policy or select **Delete** to remove obsolete NAT policies from the list of those available to a controller profile.

The screenshot shows a configuration window for a NAT pool named 'nat1'. It features a table for defining IP address ranges. The table has three columns: 'Start IP', 'End IP', and a delete icon. The first row is populated with '192.168.1.100' and '192.168.1.192'. Below the table is an 'Add Row' button. At the bottom of the window are 'OK', 'Reset', and 'Exit' buttons.

Start IP	End IP	
192.168.1.100	192.168.1.192	

FIGURE 153 NAT Pool screen

8. If adding a new NAT policy or editing the configuration of an existing policy, define the following parameters:

Name	If adding a new NAT policy, provide a name to help distinguish it from others with similar configurations. The length cannot exceed 64 characters.
Prefix Length	Use the spinner control to set the netmask (between 1 - 30) of the network the pool address belongs to.
IP Address Range	Define a range of IP addresses that are hidden from the public Internet. NAT modifies network address information in the defined IP range while in transit across a traffic routing device. NAT only provides IP address translation and does not provide a firewall. A branch deployment with NAT by itself will not block traffic from being potentially routed through a NAT device. Consequently, NAT should be deployed with a stateful firewall.

9. Select the **+ Add Row** button as needed to append additional rows to the IP Address Range table.
10. Select **OK** to save the changes or overrides made to the profile's NAT Pool configuration. Select **Reset** to revert to the last saved configuration.
11. Select the **Static NAT** tab.
The Source tab displays by default.

Source IP	NAT IP	Network	
148.98.26.48	10.1.1.2	inside	

FIGURE 154 Profile Overrides - Static NAT screen

12. To map a source IP address from an internal network to a NAT IP address click the **+ Add Row** button. Enter the internal network IP address in **Source IP** field. Enter the NAT IP address in the **NAT IP** field.
13. Use the **Network** drop-down menu to set the NAT type either *Inside* or *Outside*. Select **Inside** to create a permanent, one-to-one mapping between an address on an internal network and a perimeter or external network. To share a Web server on a perimeter interface with the Internet, use static address translation to map the actual address to a registered IP address. Static address translation hides the actual address of the server from users on insecure interfaces. Casual access by unauthorized users becomes much more difficult. Static NAT requires a dedicated address on the outside network for each host. Inside NAT is the default setting.
14. Select the **Destination** tab to view destination NAT configurations and define packets passing through the NAT on the way back to the managed LAN are searched against to the records kept by the NAT engine. The destination IP address is changed back to the specific internal private class IP address to reach the LAN over the managed network.

Static NAT creates a permanent, one-to-one mapping between an address on an internal network and a perimeter or external network. To share a Web server on a perimeter interface with the Internet, use static address translation to map the actual address to a registered IP address. Static address translation hides the actual address of the server from users on insecure interfaces. Casual access by unauthorized users becomes much more difficult. Static NAT requires a dedicated address on the outside network for each host.

Protocol	Select the protocol for use with static translation. TCP, UDP and Any are available options. TCP is a transport layer protocol used by applications requiring guaranteed delivery. It's a sliding window protocol handling both timeouts and retransmissions. TCP establishes a full duplex virtual connection between two endpoints. Each endpoint is defined by an IP address and a TCP port number. The <i>User Datagram Protocol</i> (UDP) offers only a minimal transport service, non-guaranteed datagram delivery, and provides applications direct access to the datagram service of the IP layer. UDP is used by applications not requiring the level of service of TCP or are using communications services (multicast or broadcast delivery) not available from TCP. The default setting is Any.
Destination IP	Enter the local address used at the (source) end of the static NAT configuration. This address (once translated) will not be exposed to the outside world when the translation address is used to interact with the remote destination.
Destination Port	Use the spinner control to set the local port number used at the (source) end of the static NAT configuration. The default value is port 1.
NAT IP	Enter the IP address of the matching packet to the specified value. The IP address modified can be either source or destination based on the direction specified.
NAT Port	Enter the port number of the matching packet to the specified value. This option is valid only if the direction specified is destination.
Network	Select Inside or Outside NAT as the network direction. Inside is the default setting.

17. Select **OK** to save the changes or overrides made to the static NAT configuration. Select **Reset** to revert to the last saved configuration.
18. Select the **Dynamic NAT** tab.

Dynamic NAT configurations translate the IP address of packets going out from one interface to another interface based on configured conditions. Dynamic NAT requires packets be switched through a NAT router to generate translations in the controller translation table.

NAT Pool Static NAT Dynamic NAT					
Source List ACL	Network	Interface	Overload Type	NAT Pool	Overload IP
IPACLRule1	inside	vwan1	NAT Pool	nat1	

Type to search in tables Row Count: 1

Add **Edit** **Delete** **Exit**

FIGURE 157 Profile Overrides - Dynamic NAT screen

19. Refer to the following to determine whether a new Dynamic NAT configuration requires creation, edit or deletion:

Source List ACL	Lists an ACL name to define the packet selection criteria for the NAT configuration. NAT is applied only on packets which match a rule defined in the access-list. These addresses (once translated) are not exposed to the outside world when the translation address is used to interact with the remote destination.
Network	Displays Inside or Outside NAT as the network direction for the dynamic NAT configuration.
Interface	Lists the VLAN (between 1 - 4094) used as the communication medium between the source and destination points within the NAT configuration.
Overload Type	Select the check box of Overload Type used with the listed IP ACL rule. Options include <i>NAT Pool</i> , <i>One Global Address</i> and <i>Interface IP Address</i> . Interface IP Address is the default setting.
NAT Pool	Displays the name of an existing NAT pool used with the dynamic NAT configuration.
Overload IP	If <i>One Global IP Address</i> is selected as the Overload Type, define an IP address used a filter address for the IP ACL rule.

20. Select **Add** to create a new Dynamic NAT configuration, **Edit** to modify or override an existing configuration or **Delete** to permanently remove a configuration.

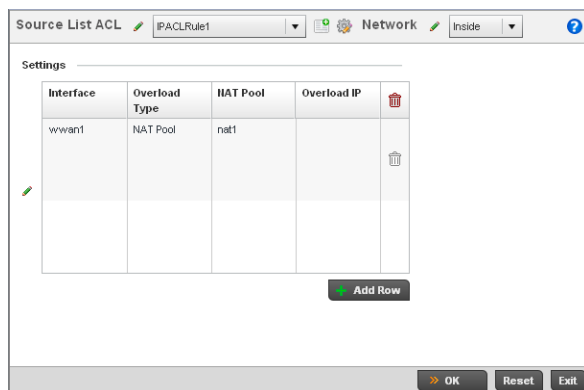


FIGURE 158 Dynamic NAT Add screen

21. Set or override the following to define the Dynamic NAT configuration:

- Source List ACL** Use the drop-down menu to select an ACL name to define the packet selection criteria for NAT. NAT is applied only on packets which match a rule defined in the access-list. These addresses (once translated) will not be exposed to the outside world when the translation address is used to interact with the remote destination.
- Network** Select *Inside* or *Outside* NAT as the network direction for the dynamic NAT configuration. Inside is the default setting.
- Interface** Use the drop-down menu to select the wireless WAN or VLAN ID (between 1 - 4094) used as the communication medium between the source and destination points within the NAT configuration. Ensure the VLAN selected represents the intended network traffic within the NAT supported configuration. VLAN1 is available by default.
- Overload Type** Select the check box of Overload Type used with the listed IP ACL rule. Options include *NAT Pool*, *One Global Address* and *Interface IP Address*. Interface IP Address is the default setting.
- NAT Pool** Provide the name of an existing NAT pool for use with the dynamic NAT configuration.
- Overload IP** If *One Global IP Address* is selected as the Overload Type, define an IP address used as a filter address for the IP ACL rule.

22. Select **OK** to save the changes or overrides made to the dynamic NAT configuration. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's Services Configuration

A controller profile can contain specific guest access (captive portal), DHCP server and RADIUS server configurations supported by the controller's own internal resources. These controller access, IP assignment and user authorization resources can be defined uniquely as controller profile requirements dictate.

To define or override a profile's services configuration:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Services**.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

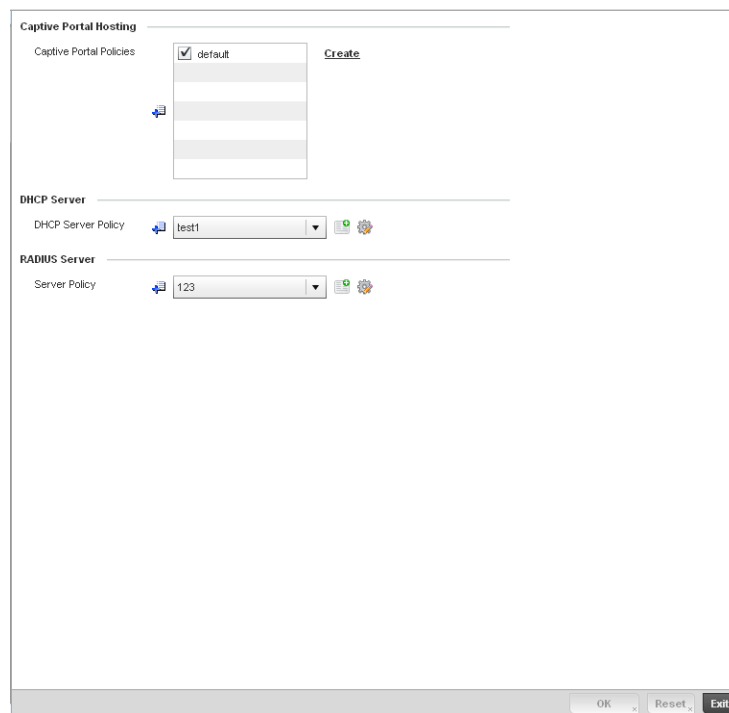


FIGURE 159 Profile Overrides - Services screen

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

5. Refer to the **Captive Portal Hosting** section to set or override a controller guest access configuration (captive portal) for use with this profile.

A *captive portal* is guest access policy for providing guests temporary and restrictive access to the managed wireless network. The primary means of securing such controller guest access is a hotspot.

A captive portal policy's hotspot configuration provides secure authenticated controller access using a standard Web browser. Hotspots provides authenticated access by capturing and re-directing a wireless user's Web browser session to a captive portal login page where the user must enter valid credentials to access to the managed wireless network. Once logged into the managed hotspot, additional *Agreement*, *Welcome* and *Fail* pages provide the administrator with a number of options on the hotspot's screen flow and user appearance.

Either select an existing captive portal policy, use the default captive portal policy or select the **Create** link to create a new captive portal configuration that can be applied to this profile. For more information, see [Configuring a Captive Portal Policy on page 10-465](#).

6. Use the **DHCP Server Policy** drop-down menu assign this controller profile a DHCP server policy. If an existing DHCP policy does not meet the profile's requirements, select the **Create** icon to create a new policy configuration that can be applied to this profile or the **Edit** icon to modify the parameters of an existing DHCP Server policy.

Dynamic Host Configuration Protocol (DHCP) allows hosts on an IP network to request and be assigned IP addresses as well as discover information about the managed network where they reside. Each subnet can be configured with its own address pool. Whenever a DHCP client requests an IP address, the DHCP server assigns an IP address from that subnet's address pool. When the controller's onboard DHCP server allocates an address for a DHCP client, the client is assigned a lease, which expires after an pre-determined interval. Before a lease expires, wireless clients (to which leases are assigned) are expected to renew them to continue to use the addresses. Once the lease expires, the client is no longer permitted to use the leased IP address. The controller profile's DHCP server policy ensures all IP addresses are unique, and no IP address is assigned to a second client while the first client's assignment is valid (its lease has not expired).

7. Either select an existing captive portal policy or select the **Create** button to create a new captive portal configuration that can be applied to this profile. For more information, see [Configuring a Captive Portal Policy on page 10-465](#).
8. Use the **RADIUS Server Policy** drop-down menu to select an existing RADIUS server policy to use as a user validation security mechanism with this controller profile.

A controller profile can have its own unique RADIUS server policy to authenticate users and authorize access to the managed network. A profile's RADIUS policy provides the centralized management of controller authentication data (usernames and passwords). When an client attempts to associate to the controller, the controller sends the authentication request to the RADIUS server. An

AP-7131 model Access Point is also equipped with its own RADIUS server.

9. If an existing RADIUS server policy does not meet your requirements, select the **Create** icon to create a new policy or the **Edit** icon to modify the parameters of an existing policy. For more information, see [Setting the Controller's RADIUS Configuration on page 10-491](#).
10. Select **OK** to save the changes or overrides made to the profile's services configuration. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's Management Configuration

The controller has mechanisms to allow/deny management access to the managed network for separate interfaces and protocols (HTTP, HTTPS, Telnet, SSH or SNMP). These management access configurations can be applied strategically to controller profiles as controller resource permissions dictate for the profile. Additionally, overrides can be applied to customize a device's management configuration, if deployment requirements change an a devices configuration must be modified from its original device profile configuration.

Additionally, an administrator can define a profile with unique configuration file and device firmware upgrade support. In a clustered environment, these operations can be performed on one controller, then propagated to each member of the cluster and onwards to devices managed by each cluster member.

To define or override a profile's management configuration:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.

4. Select **Management**.

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

FIGURE 160 Profile Overrides - Management Settings screen

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

5. Refer to the **Management Policy** field to set or override a controller management configuration for use with this profile. A default management policy is also available if no existing policies are usable.

Use the drop-down menu to select an existing management policy to apply to this controller profile. If no management policies exist meeting the data access requirements of this controller profile, select the **Create** icon to access a series of screens used to define administration, access control and SNMP configurations. Select an existing policy and select the **Edit** icon to modify the configuration of an existing management policy.
6. Use to the **Critical Resource Policy** pulldown to set or override a critical resource policy for use with this profile. For more information on defining a critical resource policy, see *Critical Resource Policy on page 5-222*.

7. Refer to the **Message Logging** section to define how the controller profile logs system events. It's important to log individual events to discern an overall pattern that may be negatively impacting controller performance using the configuration defined for this profile.

Enable Message Logging	Select the check box to enable the controller profile to log system events to a user defined log file or a syslog server. Selecting this check box enables the rest of the parameters required to define the profile's logging configuration. This option is disabled by default.
Remote Logging Host	Use this table to define numerical (non DNS) IP addresses for up to three external resources where logged system events can be sent on behalf of the controller profile. Select Clear as needed to remove an IP address.
Facility to Send Log Messages	Use the drop-down menu to specify the local server facility (if used) for the controller profile event log transfer.
Syslog Logging Level	Event severity coincides with the syslog logging level defined for the profile. Assign a numeric identifier to log events based on criticality. Severity levels include 0 - <i>Emergency</i> , 1 - <i>Alert</i> , 2 - <i>Critical</i> , 3 - <i>Errors</i> , 4 - <i>Warning</i> , 5 - <i>Notice</i> , 6 - <i>Info</i> and 7 - <i>Debug</i> . The default logging level is 4.
Console Logging Level	Event severity coincides with the syslog logging level defined for the profile. Assign a numeric identifier to log events based on criticality. Severity levels include 0 - <i>Emergency</i> , 1 - <i>Alert</i> , 2 - <i>Critical</i> , 3 - <i>Errors</i> , 4 - <i>Warning</i> , 5 - <i>Notice</i> , 6 - <i>Info</i> and 7 - <i>Debug</i> . The default logging level is 4.
Buffered Logging Level	Event severity coincides with the syslog logging level defined for the profile. Assign a numeric identifier to log events based on criticality. Severity levels include 0 - <i>Emergency</i> , 1 - <i>Alert</i> , 2 - <i>Critical</i> , 3 - <i>Errors</i> , 4 - <i>Warning</i> , 5 - <i>Notice</i> , 6 - <i>Info</i> and 7 - <i>Debug</i> . The default logging level is 4.
Time to Aggregate Repeated Messages	Define the increment (or interval) system events are logged on behalf of this controller profile. The shorter the interval, the sooner the event is logged. Either define an interval in <i>Seconds</i> (0 - 60) or <i>Minutes</i> (0 -1). The default value is 0 seconds.
Forward Logs to Controller	Select the checkbox to define a log level for forwarding event logs to the control. Log levels include <i>Emergency</i> , <i>Alert</i> , <i>Critical</i> , <i>Error</i> , <i>Warning</i> , <i>Notice</i> , <i>Info</i> and <i>Debug</i> . The default logging level is <i>Error</i> .

8. Refer to the **System Event Messages** section to define or override how controller system messages are logged and forwarded on behalf of the controller profile.

Event System Policy	Select an Event System Policy from the drop-down menu. If an appropriate policy does not exist click the create button to make a new policy.
Enable System Events	Select the Enable System Events check box to allow the controller profile to capture system events and append them to a log file. It's important to log individual events to discern an overall pattern that may be negatively impacting controller performance. This setting is enabled by default.
Enable System Event Forwarding	Select the Enable System Event Forwarding radio button to enable the forwarding of system events to another controller or cluster member. This setting is enabled by default.

9. Refer to the **Events E-mail Notification** section to define or override how system event notification e-mails are sent.

SMTP Server	Specify either the Hostname or IP Address of the outgoing SMTP server where notification e-mails will be sent from.
Port of SMTP	If a non-standard SMTP port is used on the outgoing SMTP server check this box and specify a port between 1 and 65,535 for the outgoing SMTP server to use.
Sender E-mail Address	Specify the e-mail address that notification e-mails will be sent from. This will be the from address on notification e-mails.
Username for SMTP Server	Specify the username of the sender on the outgoing SMTP server. Many SMTP servers require users to authenticate with an username and password before sending e-mail through the server.
Password for SMTP Server	Specify the password associated with the username of the sender on the outgoing SMTP server. Many SMTP servers require users to authenticate with an username and password before sending e-mail through the server.
Persist Configuration Across Reloads	Use the pull-down menu to configure whether configuration overrides should persist when the device configuration is reloaded. Available options are Enabled, Disabled and Secure.

10. Select **OK** to save the changes and overrides made to the profile's Management Settings.
Select **Reset** to revert to the last saved configuration.

11. Select **Firmware** from the Management menu.

FIGURE 161 Profile Overrides - Management Firmware screen

12. Refer to the **Auto Install via DHCP Option** section to configure automatic configuration file and firmware updates.

Enable Configuration Update	Select the Enable Configuration Update check box (from within the Automatic Configuration Update field) to enable automatic configuration file updates for the controller profile from a location external to the controller. If enabled (the setting is disabled by default), provide a complete path to the target configuration file used in the update.
Enable Firmware Update	Select this option to enable automatic controller firmware upgrades (for this controller profile) from a user defined remote location. This value is disabled by default.

13. Refer to the **Legacy Device Firmware Management** field to define or whether Brocade Mobility 650 Access Point and AP-7131 model devices can upgrade to newer firmware versions or downgrade to legacy firmware versions.

Migration Firmware from AP7131 4.x path	Provide a complete path to the target firmware used to support a legacy AP-7131 firmware update. The length of the path cannot exceed 253 characters.
Legacy AP650 Auto Upgrade	Check this box to enable automatic firmware upgrades for all legacy AP650 Access Points connected to the controller.

14. Use the parameters within the **Automatic Adopted AP Firmware Upgrade** section to define an automatic firmware upgrade from a controller based file.

Allow Controller Upgrade	Select the access point model to upgrade to a newer firmware version using its associated Virtual Controller AP's most recent firmware file for that model. This parameter is enabled by default.
Number of Concurrent Upgrades.	Use the spinner control to define the maximum number (1 - 20) of adopted APs that can receive a firmware upgrade at the same time. The default value is 10. Keep in mind that during a firmware upgrade, the AP is offline and unable to perform its normal wireless client support function until the upgrade process is complete.

15. Select **OK** to save the changes and overrides made to the profile's Management Firmware configuration. Select **Reset** to revert to the last saved configuration.

16. Select **Heartbeat** from the Management menu.

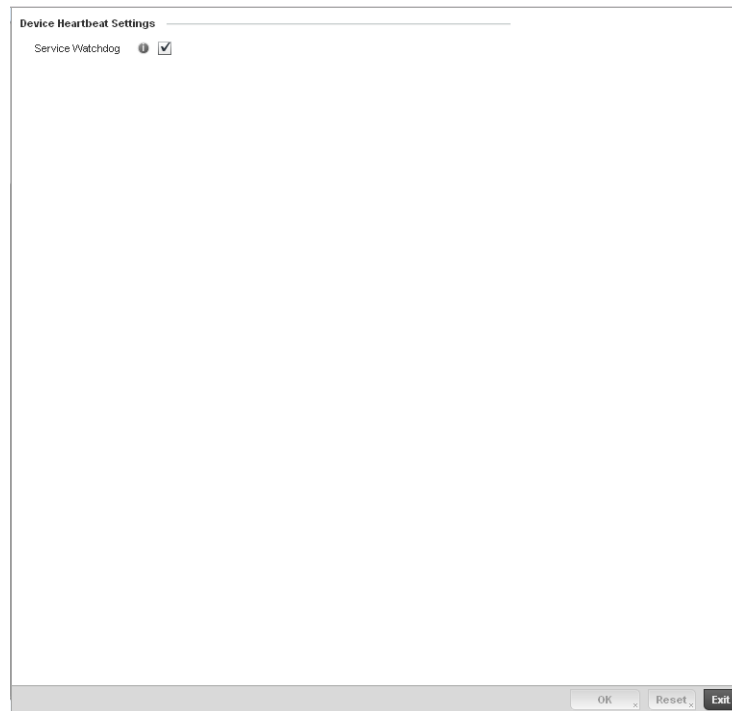


FIGURE 162 Profile Overrides - Management Heartbeat screen

17. Select the **Service Watchdog** option to implement heartbeat messages to ensure other associated devices are up and running and capable of effectively interoperating with the controller. The Service Watchdog is enabled by default.
18. Select **OK** to save the changes and overrides made to the profile maintenance Heartbeat tab. Select **Reset** to revert to the last saved configuration.

Overriding a Profile's Advanced Configuration

A profile's advanced configuration is comprised of defining its MINT protocol configuration and the profile's NAS identifier and port ID attributes. MINT provides secure controller profile communications at the transport layer. Using MINT, a device can be configured to only communicate with other authorized (MINT enabled) devices. Therefore, MINT is well designed for controller profile support, wherein a group of managed devices share the same configuration attributes. However, a profile's advanced configuration may require an override as a particular device's configuration requires a specific parameter be updated in a manner that deviates the configuration from that the configuration shared by the devices using the profile

To set or override a profile's advanced configuration:

1. Select the Devices from the Web UI.
2. Select **Profile Overrides** to expand its menu items
3. Select **Advanced** to expand its sub menu items.

MINT provides the means to secure controller profile communications at the transport layer. Using MINT, a device can be configured to only communicate with other authorized (MINT enabled) devices.

managed devices can communicate with each other exclusively over a MINT security domain. Keys can also be generated externally using any application (like openssl). These keys must be present on the managed device managing the domain for key signing to be integrated with the UI. A MAP device that needs to communicate with another first negotiates a security context with that device.

The security context contains the transient keys used for encryption and authentication. A secure network requires users to know about certificates and PKI. However, administrators do not need to define security parameters for Access Points to be adopted (secure WISPe being an exception, but that isn't a commonly used feature). Also, users can replace any device on the network or move devices around and they continue to work. Default security parameters for MINT are such that these scenarios continue to function as expected, with minimal user intervention required only when a new network is deployed.

To define or override a controller profile's MINT configuration:

4. Select **MINT Protocol** from the Advanced menu item.

The screenshot shows a configuration window titled "Settings" with tabs for "Settings", "IP", and "VLAN". The "Settings" tab is active. The configuration is organized into sections:

- Area Identifier:** "Level 1 Area ID" is set to 1, with a checkbox for a spinner control checked. The range is (1 to 16,777,215).
- Priority Adjustment:** "Designated IS Priority Adjustment" is set to 66, with a range of (-255 to 255).
- Shortest Path First (SPF):** "Latency of Routing Recalculation" is set to 17, with a range of (0 to 60 seconds) and a checked checkbox for a spinner control.
- MINT Link Settings:** "MLCP IP" and "MLCP VLAN" both have checked checkboxes.

Buttons for "OK", "Reset", and "Exit" are located at the bottom right of the window.

FIGURE 163 Advanced Profile Overrides MINT screen - Settings tab

5. The **Settings** tab displays by default.
6. Refer to the **Area Identifier** field to define or override the Level 1 and Level 2 Area IDs used by the profile's MINT configuration.

Level 1 Area ID

Select the box to enable a spinner control for setting the Level 1 Area ID between 1 - 16,777,215. The default value is disabled.

12. The IP tab displays the **IP address**, **Routing Level**, **Listening Link**, **Port**, **Forced Link**, **Link Cost**, **Hello Packet Interval** and **Adjacency Hold Time** managed devices use to securely communicate amongst one another. Select **Add** to create a new Link IP configuration or **Edit** to override an existing MINT configuration.

FIGURE 165 Advanced Profile MINT screen - IP tab

13. Set the following **Link IP** parameters to complete the MINT network address configuration:

IP	Define or override the IP address used by peer controllers for interoperation when supporting the MINT protocol.
Port	To specify a custom port for MiNT links, check this box and use the spinner control to define or override the port number between 1 and 65,535.
Routing Level	Use the spinner control to define or override a routing level of either 1 or 2.
Listening Link	Specify a listening link of either 0 or 1. UDP/IP links can be created by configuring a matching pair of links, one on each end point. However, that is error prone and doesn't scale. So UDP/IP links can also listen (in the TCP sense), and dynamically create connected UDP/IP links when contacted. The typical configuration is for the controller to have a listening UDP/IP link on the switch's IP address S.S.S.S, and for all the APs to have a regular UDP/IP link to S.S.S.S.
Forced Link	Check this box to specify the MiNT link as a forced link.
Link Cost	Use the spinner control to define or override a link cost between 1 - 10,000. The default value is 100.
Hello Packet Interval	Set or override an interval in either <i>Seconds</i> (1 - 120) or <i>Minutes</i> (1 - 2) for the transmission of hello packets. The default interval is 15 seconds.
Adjacency Hold Time	Set or override a hold time interval in either <i>Seconds</i> (2 - 600) or <i>Minutes</i> (1 - 10) for the transmission of hello packets. The default interval is 46 seconds.

14. Select the **VLAN** tab to display the link IP VLAN information shared by the devices managed by the controller's MINT configuration.

Link Cost	Use the spinner control to define or override a link cost between 1 - 10,000. The default value is 100.
Hello Packet Interval	Set or override an interval in either <i>Seconds</i> (1 - 120) or <i>Minutes</i> (1 - 2) for the transmission of hello packets. The default interval is 15 seconds.
Adjacency Hold Time	Set or override a hold time interval in either <i>Seconds</i> (2 - 600) or <i>Minutes</i> (1 - 10) for the transmission of hello packets. The default interval is 46 seconds.

17. Select **OK** to save the updates and overrides to the MINT Protocol configuration. Select **Reset** to revert to the last saved configuration.

Advanced Profile Miscellaneous Configuration

Overriding a Profile's Advanced Configuration

Refer to the advanced profile's Miscellaneous menu item to set or override a profile's NAS configuration. The profile database on the RADIUS server consists of user profiles for each connected *network access server* (NAS) port. Each profile is matched to a username representing a physical port. When the wireless controller authorizes users, it queries the user profile database using a username representative of the physical NAS port making the connection. Access Point LED behavior and RF Domain management can also be defined from within the Miscellaneous screen.

1. Select **Miscellaneous** from the Advanced menu item

The screenshot shows a configuration window titled "Device RADIUS Authentication Parameters". It has two input fields for "NAS-Identifier Attribute" and "NAS-Port-Id Attribute". Below these is a section for "RF Domain Manager" with a checked "Capable" checkbox and a "Priority" spinner set to 1. The bottom right corner contains "OK", "Reset", and "Exit" buttons.

FIGURE 168 Advanced Profile Overrides - Miscellaneous screen

2. Set a **NAS-Identifier Attribute** up to 253 characters in length.

This is the RADIUS NAS-Identifier attribute that typically identifies the Access Point or controller of controller where a RADIUS message originates.

3. Set a **NAS-Port-Id Attribute** up to 253 characters in length.
This is the RADIUS NAS port ID attribute which identifies the device port where a RADIUS message originates.
4. Select the **Capable** check box (within the **RF Domain Manager** section) to designate this specific profile managed device as being capable of being the RF Domain manager for a particular RF Domain. The default value is enabled.
5. Select the **Priority** check box (within the **RF Domain Manager** section) to set a priority value for this specific profile managed device. Once enabled, use the spinner control to set a device priority between 1 - 255. The higher the number set, the higher the priority in the RF Domain manager election process.
6. Select **OK** to save the changes made to the profile's Advanced Miscellaneous configuration. Select **Reset** to revert to the last saved configuration.

Auto Provisioning Policies

Wireless devices running Brocade Mobility can adopt other wireless devices. For example, a wireless controller can adopt an number of access points. When a device is adopted the device configuration is determined by the adopting device. Since multiple configuration policies are supported an adopting device needs a way of determining which of the multiple configuration policies should be used for a given adoptee. Auto Provisioning Policies provide a way to determine a configuration policy to be used for an adoptee based on some of its properties. For example, a configuration policy could be assigned based on a MAC address, IP address, CDP snoop strings, etc.

Once created an auto provisioning policy can be used in profiles or device configuration objects. An auto provisioning policy contains a set of ordered by precedence rules that either deny or allow adoption based on a potential adoptee properties and a catch-all variable that determines if the adoption should be allowed when none of the rules were matched. All rules (both deny and allow) are evaluated sequentially starting with the rule with the lowest precedence value. The evaluation stops as soon as a rule has been matched, no attempt is made to find a better match further down in the set.

The evaluation is performed using various matching criteria. The matching criteria supported in Brocade Mobility are:

MAC	Matches the MAC address of a device attempting to be adopted. Either a single MAC address or a range of MAC addresses can be specified.
VLAN	Matches when adoption over a Layer 2 link matches the VLAN ID of an adoption request. Note that this is a VLAN ID as seen by the recipient of the request, in case of multiple hops over different VLANs this may differ from VLAN ID set by the sender. A single VLAN ID is specified in the rule. This rule is ignored for adoption attempts over Layer 3.
IP Address	Matches when adoption is using a Layer 3 link matches the source IP address of an adoption request. In case of NAT the IP address may be different from what the sender has used. A single IP, IP range or IP/mask is specified in the rule. This rule is ignored for adoption attempts over Layer 2.
Serial Number	Matches exact serial number (case insensitive).
Model	Matches exact model name (case insensitive).

DHCP Option	Matches the value found in DHCP vendor option 191 (case insensitive). DHCP vendor option 191 can be setup to communicate various configuration parameters to an AP. The value of the option is a string in the form of tag=value separated by a semicolon, e.g. 'tag1=value1;tag2=value2;tag3=value3'. The access point includes the value of tag 'rf-domain', if present. This value is matched against the auto provisioning policy.
FQDN	Matches a substring to Fully Qualified Domain Name of a device (case insensitive).
CDP	Matches a substring in a list of CDP snoop strings (case insensitive). For example, if an access point snooped 3 devices: controller1.moto.com, controller2.moto.com and controller3.moto.com, 'controller1', 'moto', 'moto.com', are examples of the substrings that will match.
LLDP	Matches a substring in a list of LLDP snoop strings (case insensitive). For example, if an access point snooped 3 devices: controller1.moto.com, controller2.moto.com and controller3.moto.com, 'controller1', 'moto', 'moto.com', are examples of the substrings that will match.

Auto Provisioning is the process an Access Point uses to discover controllers available in the network, pick the most desirable controller, establish an association, optionally obtain an image upgrade and obtain its configuration.

At adoption, an Access Point solicits and receives multiple adoption responses from controllers available on the network. These adoption responses contain loading policy information the Access Point uses to select the optimum controller for adoption. By default, an auto provisioning policy generally distributes AP adoption evenly amongst available controllers. Modify existing adoption policies or create a new one as needed to meet the adoption requirements of a device and their assigned controller profile.

NOTE

A device configuration does not need to be present for an auto provisioning policy to take effect. Once adopted, and the device's configuration is defined and applied by the controller, the auto provisioning policy mapping does not have impact on subsequent adoptions by the same device.

An auto provisioning policy enables an administrator to define adoption rules for the supported Brocade Mobility Access Points capable of adoption by a wireless controller.

Auto provisioning policies set the different restrictions on how an AP gets adopted to a wireless controller managed network.

To review existing Auto Provisioning Policy configurations:

1. Select **Configuration > Devices > Auto Provisioning Policy**.
2. The **Adoption** screen displays by default.

- Argument 2** The number of arguments vary on the Match Type. This column lists the second argument value. This value is not set as part of the rule creation or edit process.
- RF Domain Name** Sets the name of the RF Domain to which the device is adopted automatically. Select the **Create** icon to define a new RF Domain configuration or select the **Edit** icon to revise an existing configuration. For more information, see to *Managing RF Domains on page 8-408*.
- Profile Name** Defines the name of the profile used when the Auto Provisioning Policy is applied to a device. Select the **Create** icon to define a new Profile configuration or select the **Edit** icon to revise an existing configuration. For more information, see [General Profile Configuration on page 7-316](#).
5. If a rule requires addition or modification, select either **Add** or **Edit** to define the required parameters using the **Rule** screen.

FIGURE 171 Auto Provisioning Policy Rule screen

6. Specify the following parameters in the **Rule** screen:

- Rule Precedence** Use the spinner control to specify the precedence (sequence) that the Adoption Policies rules are applied. Rules with the lowest precedence receive the highest priority. This value is set (between 1 - 1000) when adding a new Auto Provisioning Policy rule configuration.
- Device Type** Set the Brocade Mobility 650 Access Point, AP7131, or AP6511 Access Point model for which this policy applies. Adoption rules are specific to the selected model.

- Match Type** Set the matching criteria used in the policy. This is like a filter and further refines the APs that can be adopted. The **Match Type** can be one of the following:
MAC Address – The filter type is a MAC Address of the selected Access Point model.
IP Address – The filter type is the IP address of the selected Access Point model.
VLAN – The filter type is a VLAN.
Serial Number – The filter type is the serial number of the selected Access Point model.
Model Number – The filter type is the Access Point model number.
DHCP Option – The filter type is the DHCP option value of the selected Access Point model.
- RF Domain Name** Set the name of the RF Domain to which the device is adopted automatically. Select the **Create** icon to define a new RF Domain configuration or select the **Edit** icon to revise an existing configuration. For more information, see to *General Profile Configuration on page 7-316*.
- Profile Name** Define the name of the profile used when the Auto Provisioning Policy is applied to a device. Select the **Create** icon to define a new Profile configuration or select the **Edit** icon to revise an existing configuration. For more information, see [General Profile Configuration on page 7-316](#).

7. Select the **Default** tab to define the Auto Provisioning Policy's rule matching adoption configuration.




FIGURE 172 Auto Provisioning Policy screen - Default tab

8. Select the **Adopt if No Rules Match** checkbox to have the controller adopt when no matching filter rules apply. This setting is enabled by default.
9. Select **OK** to save the updates to the Auto Provisioning Policy screen. Selecting **Reset** reverts the screen to the last saved configuration.

Critical Resource Policy

A Critical Resource Policy defines a list of device IP addresses on the network (gateways, routers etc.). The support of these defined IP address is interpreted as critical to the health of the managed network. These devices addresses are pinged regularly by the wireless controller. If there's a connectivity issue, an event is generated stating a critical resource is unavailable. There's no restoration of the critical device involved.

To define a Critical Resource Policy:

1. Select **Configuration > Devices > Critical Resource Policies**.

Critical Resource Policy Name	Ping Interval
eng	30s

FIGURE 173 Critical Resource Policy screen

2. Refer to the following to help determine whether a new Critical Resource Policy should be created or an existing policy modified:

Critical Resource Policy Name Displays the name of the policy assigned when it was initially created. The policy is a collection of critical resources grouped logically.

Ping Interval The interval between 2 pings to the critical resource. Ping is used to check if connection to the critical resource is working.

1. Select **Add** to create a new policy or **Edit** to modify an existing Critical Resource Policy configuration. For more information, refer to *Managing Critical Resource Policies* on page 5-223.

Managing Critical Resource Policies

Critical Resource Policy

The controller provides some flexibility to define new IP addresses interpreted as critical resources or remove addresses no longer defined as critical.

To add or modify a Critical Resource Policy:

1. Select **Add** or **Edit** (after selecting an existing policy) from the Critical Resource Policy screen.
2. If adding a new policy, enter a name in the **Critical Resource Policy** field. Click the **OK** button (which flashes after inputting a policy name) to fill in the rest of the information for creating a Critical Resource Policy. The following screen displays.

Ping Interval

Ping Interval Seconds (5 to 86,400)

Critical Resource List

IP Address	Ping Mode	VLAN	
192.68.181.3	arp-only	1	

FIGURE 174 Critical Resource Policy Configuration screen

3. Set the following Critical Resource Policy parameters:

Ping Interval	Set the duration between two successive pings to the critical device. Select from: <i>Days</i> – Measured in days. <i>Hours</i> – Measured in hours. <i>Minutes</i> – Measured in minutes <i>Seconds</i> – Measured in seconds The default interval is 30 seconds.
IP Address	Set the IP address of the critical resource. This is the address the device is assigned and is used by the wireless controller to check if the critical resource is available.
Ping Mode	Set the ping mode used when the availability of a critical resource is validated. Select from: <i>arp-only</i> – Use the <i>Address Resolution Protocol (ARP)</i> only for pinging the critical resource. ARP is used to resolve hardware addresses when only the network layer address is known. <i>arp-icmp</i> – Use both <i>Address Resolution Protocol (ARP)</i> and <i>Internet Control Message Protocol (ICMP)</i> for pinging the critical resource and sending the control messages (device not reachable, requested service not available, etc).
VLAN	Define the VLAN on which the critical resource is available. Enter the VLAN number in the text provided or select the VLAN using the spinner control.

4. Click the **Add Row** button at the bottom of the **Critical Resource List** table to add a new critical resource. To edit an existing critical resource, select the row and edit the values.
5. Select **OK** to save the changes. Select **Reset** to revert to the last saved configuration. Delete obsolete rows as needed.

Managing Event Policies

Critical Resource Policy

Event Policies enable an administrator to create specific notification mechanisms using one, some or all of the SNMP, syslog, controller forwarding or email notification options available to the controller. Each listed event can have customized notification settings defined and saved as part of an event policy. Thus, policies can be configured and administrated in respect to specific sets of client association, authentication/encryption and performance events. Once policies are defined, they can be mapped to device profiles strategically as the likelihood of an event applies to particular devices.

When initially displayed, the Event Policy screen lists existing policies. Existing policies can have their event notification configurations modified as device profile requirements warrant. New policies can be added as needed.

To add or modify an Event Policy:

1. Select **Add** or **Edit** (after selecting an existing policy) from the Event Policy screen.
2. If adding a new policy, enter a name in the **Policy Name** field. Click the **OK** button (which flashes after inputting a policy name) to fill in the rest of the information for creating a Event Policy. The following screen displays.

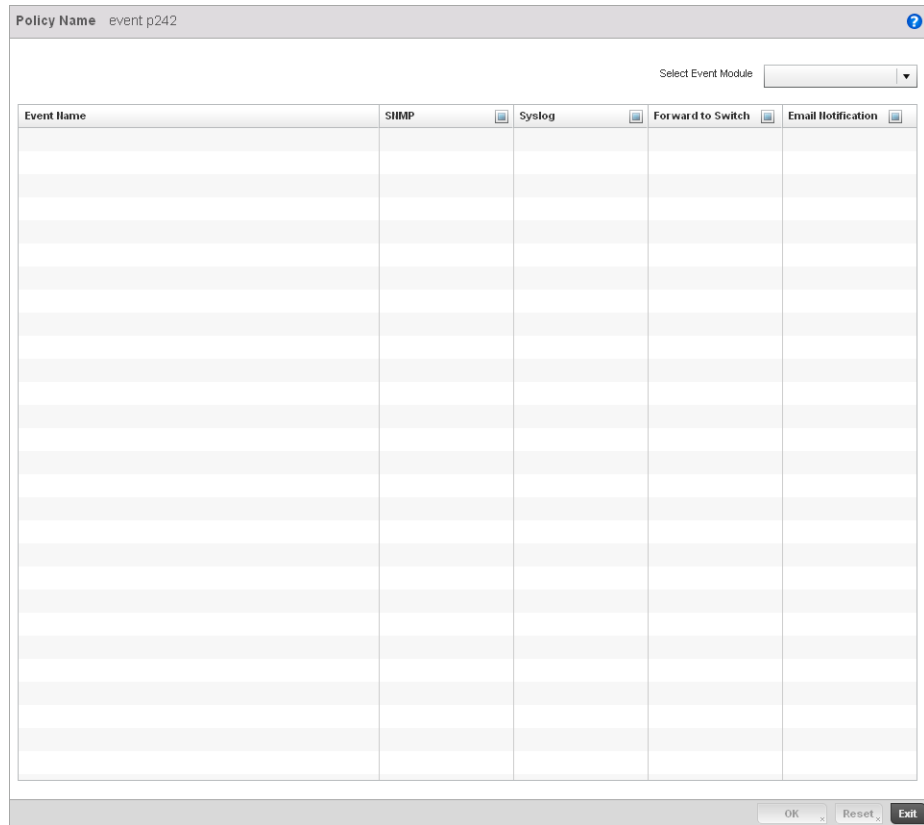


FIGURE 175 Event Policy Configuration screen

3. Refer to the **Select Event Module** drop-down menu on the top right-hand side of the screen, and select controller event module used to track the occurrence of each list event.

Review each event and select (or deselect) the **SNMP**, **Syslog**, **Forward to Switch** or **Email Notification** option as required for the event and applicable device profile. Map an existing policy to a device profile as needed. Select Profile from the Map drop-down menu in the lower-left hand side of the screen. Expand the list of device profiles available, and apply the event policy as required.

4. Select **OK** to save the changes. Select **Reset** to revert to the last saved configuration. Delete obsolete rows as needed.

Managing MINT Policies

Critical Resource Policy

To add or modify a MINT Policy:

1. Select **Devices** > **MINT Policy** to display the MINT Policy screen.

The screenshot shows a configuration window titled "Settings" with the following parameters:

- Level 2 Area ID:** A numeric input field containing "1", with a range of "(1 to 16,777,215)".
- MTU:** A numeric input field containing "1500", with a range of "(900 to 1,500)".
- UDP/IP Encapsulation Port:** A numeric input field containing "24576", with a range of "(2 to 65,534)".

At the bottom right of the window, there are three buttons: "OK", "Reset", and "Exit".

FIGURE 176 MINT Policy Configuration screen

2. Configure the following parameters to configure the MINT policy:

Level 2 Area ID	Define a Level 2 Area ID for the Mint Policy. The Level 2 Area ID is the global mint area identifier. This area identifier separates two overlapping mint networks and need only be configured if the administrator has two mint networks that share the same packet broadcast domain.
MTU	Specify a MTU value for the mint policy between 900 and 1,500. The MTU setting specifies the maximum packet size that will be used for mint packets. Larger packets will be fragmented so they fit within this packet size limit. The administrator may want to configure this parameter if the mint backhaul network requires or recommends smaller packet sizes. The default value is 1500.
UDP/IP Encapsulation Port	Specify the port to use for UDP/IP encapsulation between 2 and 65,534. This value specifies an alternate UDP port to be used by mint packets and must be an even number. This port number will be used by mint control packets, and this port value plus 1 will be used to carry mint data packets. The default value is 24576.

3. Select **OK** to save the changes. Select **Reset** to revert to the last saved configuration.

Wireless Configuration

In this chapter

- [Wireless LAN Policy](#) 228
- [WPA-TKIP Deployment Considerations](#) 242
- [Configuring WLAN QoS Policies](#) 264
- [Radio QoS Policy](#) 278
- [AAA Policy](#) 289
- [Association ACL](#) 298
- [Smart RF Policy](#) 301

A *Wireless Local Area Network* (WLAN) is a data-communications system and wireless local area network that flexibly extends the functionalities of a wired LAN. A WLAN links two or more computers or devices using spread-spectrum or OFDM modulation based technology. A wireless controller managed WLAN does not require lining up devices for line-of-sight transmission, and are thus, desirable for wireless controller managed wireless networking. Roaming users can be handed off from one wireless controller connected access point to another, like a cellular phone system. WLANs can therefore be configured around the needs of specific user groups, even when they are not in physical proximity.

WLANs can be used to provide an abundance of services, including data communications (allowing mobile devices to access applications), email, file and print services or even specialty applications (such as guest access control and asset tracking).

Each wireless controller WLAN configuration contains encryption, authentication and QoS policies and conditions for user connections. Connected access point radios transmit periodic beacons for each BSS. A beacon advertises the SSID, security requirements, supported data rates of the wireless network to enable clients to locate and connect to the wireless controller managed WLAN.

WLANs are mapped to radios on each connected Brocade Mobility 650 Access Point, AP6511, or AP-7131 (adaptive mode) access point. A WLAN can be advertised from a single access point radio or can span multiple access points and radios. WLAN configurations can be defined to only provide service to specific areas of a site. For example a guest access WLAN may only be mapped to a 2.4GHz radio in a lobby or conference room providing limited coverage while a data WLAN is mapped to all 2.4GHz and 5GHz radios at the branch site providing complete coverage.

Brocade Mobility RFS4000 and RFS6000 model wireless controllers support a maximum of 32 WLANs. The Brocade Mobility RFS7000 model wireless controller supports up to 256 WLANs.

The controller's wireless configuration is comprised the following policies:

- [Wireless LAN Policy](#)
- [Configuring WLAN QoS Policies](#)
- [Radio QoS Policy](#)
- [AAA Policy](#)

- [Association ACL](#)
- [Smart RF Policy](#)

These parameters can be separately selected within the *Configuration > Wireless* pane located in top, left-hand, side of the controller UI.

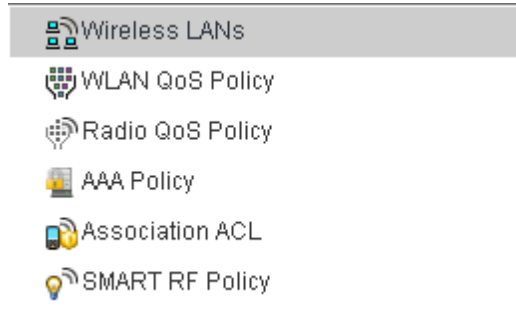


FIGURE 177 Configuration > Wireless pane

Wireless LAN Policy

To review the attributes of existing controller WLANs and, if necessary, modify their configurations:

1. Select **Configuration > Wireless > Wireless LANs** to display a high-level display of the existing WLANs.

WLAN	SSID	Description	WLAN Status	VLAN Pool	Bridging Mode	Authentication Type	Encryption Type	QoS Policy	Association ACL
WLAN1	WLAN1ESSID		Enabled	11,12,13,14	Tunnel	EAP	WEP128	WLANQoS2	
WLAN10	WLAN10ESSID		Enabled		Tunnel	None	None		
WLAN100	WLAN100ESSID		Enabled		Tunnel	None	None		
WLAN101	WLAN101ESSID		Enabled		Tunnel	None	None		
WLAN102	WLAN102ESSID		Enabled		Tunnel	None	None		
WLAN103	WLAN103ESSID		Enabled		Tunnel	None	None		
WLAN104	WLAN104ESSID		Enabled		Tunnel	None	None		
WLAN105	WLAN105ESSID		Enabled		Tunnel	None	None		
WLAN106	WLAN106ESSID		Enabled		Tunnel	None	None		
WLAN107	WLAN107ESSID		Enabled		Tunnel	None	None		
WLAN108	WLAN108ESSID		Enabled		Tunnel	None	None		
WLAN109	WLAN109ESSID		Enabled		Tunnel	None	None		
WLAN11	WLAN11ESSID		Enabled		Tunnel	None	None		
WLAN110	WLAN110ESSID		Enabled		Tunnel	None	None		
WLAN111	WLAN111ESSID		Enabled		Tunnel	None	None		
WLAN112	WLAN112ESSID		Enabled		Tunnel	None	None		
WLAN113	WLAN113ESSID		Enabled		Tunnel	None	None		
WLAN114	WLAN114ESSID		Enabled		Tunnel	None	None		
WLAN115	WLAN115ESSID		Enabled		Tunnel	None	None		
WLAN116	WLAN116ESSID		Enabled		Tunnel	None	None		
WLAN117	WLAN117ESSID		Enabled		Tunnel	None	None		
WLAN118	WLAN118ESSID		Enabled		Tunnel	None	None		
WLAN119	WLAN119ESSID		Enabled		Tunnel	None	None		
WLAN12	WLAN12ESSID		Enabled		Tunnel	None	None		
WLAN120	WLAN120ESSID		Enabled		Tunnel	None	None		
WLAN121	WLAN121ESSID		Enabled		Tunnel	None	None		
WLAN122	WLAN122ESSID		Enabled		Tunnel	None	None		
WLAN123	WLAN123ESSID		Enabled		Tunnel	None	None		
WLAN124	WLAN124ESSID		Enabled		Tunnel	None	None		

Type to search in tables

Row Count: 204

Add Edit Delete

FIGURE 178 Wireless LANs screen

2. Refer to the following (read only) information to assess the attributes of the each WLAN available to the wireless controller:

WLAN	Displays the name of each WLAN available on the wireless controller. Each WLAN can be selected and its SSID and client management properties modified.
SSID	Displays the name of the SSID assigned to the WLAN when it was created or last modified. Optionally, select a WLAN and click the Edit button to update the SSID.
Description	Displays the brief description defined for each listed WLAN when it was either created or modified.
WLAN Status	Lists each WLAN's current status as either Active or Shutdown . A green checkmark defines the WLAN as available to clients on all radios where it has been mapped. A red "X" defines the WLAN as shutdown, meaning even if the WLAN is mapped to radios, it's not available for clients to associate.
VLAN Pool	Lists each WLANs current VLAN mapping. The wireless controller permits mapping a WLAN to more than one VLANs. When a client associates with a WLAN, the client is assigned a VLAN by means of load balance distribution. The VLAN is picked from a pool assigned to the WLAN. Keep in mind however, typical deployments only map a single VLAN to a WLAN. The use of a pool is strictly optional.
Authentication Type	Displays the name of the authentication scheme this WLAN is using to secure its client membership transmissions. None is listed if authentication is not used within this WLAN. Refer to the Encryption type column if no authentication is used to verify there is some sort of data protection used with the WLAN or risk using this WLAN with no protection at all.
Encryption Type	Displays the name of the encryption scheme this WLAN is using to secure its client membership transmissions. None is listed if encryption is not used within this WLAN. Refer to the Authentication type column if no encryption is used to verify there is some sort of data protection used with the WLAN or risk using this WLAN with no protection at all.
QoS Policy	Lists the QoS policy applied to each listed WLAN. A QoS policy needs to be custom selected (or created) for each WLAN in respect to the WLAN's intended client traffic and the voice, video or normal data traffic it supports.
Association ACL	Lists the Association ACL policy applied to each listed WLAN. An Association ACL is a policy-based <i>Access Control List</i> (ACL) that either prevents or allows wireless clients from connecting to a managed WLAN. The mapping of an Association ACL is strictly optional.

Use the wireless controller's sequential set of WLAN screens to define a unique configuration for each WLAN. Refer to the following to set WLAN configurations:

- [Basic WLAN Configuration](#)
- [Configuring WLAN Security](#)
- [Configuring WLAN Firewall Support](#)
- [Configuring Client Settings](#)
- [Configuring WLAN Accounting Settings](#)
- [Configuring Client Load Balancing Settings](#)
- [Configuring Advanced WLAN Settings](#)

Basic WLAN Configuration

[Wireless LAN Policy](#)

When creating or modifying a wireless controller WLAN, the first screen that displays as part of the WLAN configuration screen flow is the Basic Configuration screen. Use this screen to enable a WLAN and define its SSID, client behavior and VLAN assignments.

1. Select **Configuration > Wireless > Wireless LAN Policy** to display a high-level display of the existing WLANs available to the wireless controller managed network.
2. Select the **Add** button to create an additional WLAN, or select an existing WLAN then **Edit** to modify the properties of the existing WLAN.

Brocade Mobility RFS4000 and RFS6000 model wireless controllers support a maximum of 32 WLANs. The Brocade Mobility RFS7000 model wireless controller supports up to 256 WLANs.

FIGURE 179 WLAN Policy Basic Configuration screen

3. Refer to the **WLAN Configuration** field to define the following:

WLAN

If adding a new WLAN, enter its name in the space provided. Spaces between words are not permitted. The name could be a logical representation of the WLAN coverage area (engineering, marketing etc.). If editing an existing WLAN, the WLAN's name appears at the top of the screen and cannot be modified. The name cannot exceed 32 characters.

SSID

Enter or modify the *Services Set Identification* (SSID) associated with the WLAN. The maximum number of characters that can be used for the SSID is 32.

Description	Provide a textual description for the WLAN to help differentiate it from others with similar configurations. The description can be up to 64 characters.
WLAN Status	Select the Enabled radio button to make this WLAN active and available to clients on all radios where it has been mapped. Select the Disabled radio button to make this WLAN inactive, meaning even if the WLAN is mapped to radios, it's not available for clients to associate and use.
QoS Policy	Use the drop-down menu to assign an existing QoS policy to the WLAN or select the Create icon to define a new QoS policy or select the Edit icon to modify the configuration of the selected QoS Policy. QoS helps ensure each WLAN receives a fair share of the controller's overall bandwidth, either equally or per the proportion configured. For information on creating a QoS policy that can be applied to WLAN, see <i>Configuring WLAN QoS Policies on page 6-264</i> .
Bridging Policy	Use the pull-down menu to specify a bridging policy for the WLAN. Available bridging policy modes are Local, Tunnel or split-tunnel.

4. Refer to the **Other Settings** field to define broadcast behavior within this specific wireless controller managed WLAN.

Broadcast SSID	Select this check box to enable the wireless controller to broadcast SSIDs within beacons. If a hacker tries to isolate and hack a client SSID via a client, the ESSID will display since the ESSID is in the beacon. This feature is enabled by default.
Answer Broadcast Probes	Select this check box to associate a client with a blank SSID (regardless of which SSID the wireless controller is currently using). This feature is enabled by default.

5. Refer to the **VLAN Assignment** field to add or remove VLANs for the selected WLAN, and define the number of clients permitted. Remember, users belonging to separate VLANs can share the same WLAN. It's not necessary to create a new WLAN for every VLAN in the network.

Single VLAN	Select the Single VLAN radio button to assign just one VLAN to this WLAN. Enter the name of the VLAN within the VLAN parameter field that displays when the Single VLAN radio button is selected. Utilizing a single VLAN per WLAN is a more typical deployment scenario than using a VLAN pool.
VLAN Pool	Select the VLAN Pool radio button to display a table with VLAN and wireless client columns (representing configurable options). Define the VLANs available to this WLAN. Additionally, define the number of wireless clients supported by each VLAN. Use the radio button's on the left-hand side of the table to enable or disable each VLAN and wireless client configuration for the WLAN. Select the + Add button to add additional VLANs to the WLAN.

6. Select the **Allow Radius Override** check box in the RADIUS VLAN Assignment to allow an Access Point to override the WLAN configuration based VLAN assigned to a wireless client and use the VLAN assigned by a RADIUS Server. If, as part of the authentication process, the RADIUS server returns a client's VLAN-ID in a RADIUS Access-Accept packet, and this feature is enabled, all client traffic is forward on that VLAN. If disabled, the RADIUS server returned VLAN-ID is ignored and the VLAN configuration (defined above) is used.
7. Select **OK** when completed to update the WLAN's basic configuration. Select **Reset** to revert the screen back to the last saved configuration.

WLAN Basic Configuration Deployment Considerations

[Basic WLAN Configuration](#)

Before defining a WLAN's basic configuration, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- Brocade Mobility recommends one VLAN be deployed for secure WLANs, while separate VLANs be defined for each WLAN using a legacy encryption scheme or providing guest access.

Configuring WLAN Security

Wireless LAN Policy

A managed WLAN can be assigned a security policy supporting authentication, captive portal (hotspot) or encryption schemes.

The controller supports a security screen where each available security option can be defined from one central location.

The screenshot displays the WLAN Policy Security configuration interface. It is organized into several sections:

- Select Authentication:** Includes radio buttons for Authentication Type (EAP, EAP-PSK, EAP-MAC, MAC, Kerberos, PSK / None). Kerberos Configuration has a Settings link. AAA Policy is a dropdown menu. Reauthentication is a checkbox with a value of 30 (range 30 to 86,400).
- Captive Portal:** Enforcement includes 'Captive Portal Enable' (checked) and 'Captive Portal if Primary Authentication Fails'. Captive Portal Policy is a dropdown menu set to 'default'.
- Select Encryption:** Includes checkboxes for WPA/WPA2-TKIP, WPA2-CCMP (checked), WEP 128, KeyGuard, WEP 64, and Open.
- Key Settings:** Pre-Shared Key field with a note 'Enter 64 HEX or 8-63 ASCII Characters' and an ASCII dropdown.
- Key Rotation:** Unicast Rotation Interval and Broadcast Rotation Interval, both set to 30 (range 30 to 86,400 seconds).
- Advanced:** TKIP Countermeasure Hold Time set to 1 Minutes (range 0 to 1,092). Exclude WPA2 TKIP is a checkbox.

Buttons for OK, Reset, and Exit are located at the bottom right.

FIGURE 180 WLAN Policy Security screen

Authentication ensures only known and trusted users or devices access a wireless controller managed WLAN. Authentication is enabled per WLAN to verify the identity of both users and devices. Authentication is a challenge and response procedure for validating user credentials such as username, password and sometimes secret-key information.

A client must authenticate to an Access Point to receive resources from the wireless controller managed network. The wireless controller supports *EAP*, *EAP PSK*, *EAP-MAC*, *Kerberos*, *MAC* and *PSK/None* authentication options.

Refer to the following to configure an authentication scheme for a wireless controller managed WLAN:

- [802.1x EAP, EAP PSK and EAP MAC](#)

- [MAC Authentication](#)
- [Kerberos](#)
- [PSK / None](#)

Secure guest access to the managed network is referred to as *captive portal*. A captive portal is guest access policy for providing guests temporary and restrictive access to the managed wireless network. The primary means of securing such controller guest access is the use of a hotspot. Existing captive portal policies can be applied to a WLAN to provide secure guest access.

A captive portal policy's hotspot configuration provides secure authenticated controller access using a standard Web browser. Hotspots provides authenticated access by capturing and re-directing a wireless user's Web browser session to a captive portal login page where the user must enter valid credentials to access to the managed wireless network. Once logged into the managed hotspot, additional Agreement, Welcome and Fail pages provide the administrator with a number of options on the hotspot's screen flow and user appearance.

Refer to *Keyguard* on page 6-249 for information on assigning a captive portal policy to a WLAN. A captive portal is a guest access configuration policy that can applied to a WLAN to provide strategic access to the WLAN and managed network.

Encryption is central for WLAN security, as it provides data privacy for traffic forwarded over a wireless controller managed WLAN. When the 802.11 specification was introduced, *Wired Equivalent Privacy* (WEP) was the primary encryption mechanism. WEP has since been interpreted as flawed in many ways, and is not considered an effective standalone encryption scheme for securing a wireless controller WLAN. WEP is typically used WLAN deployments designed to support legacy clients. New device deployments should use either WPA or WPA2 encryption.

Encryption applies a specific algorithm to alter its appearance and prevent unauthorized hacking. Decryption applies the algorithm in reverse, to restore the data to its original form. A sender and receiver must employ the same encryption/decryption method to interoperate. When both TKIP and CCMP are both enabled a mix of clients are allowed to associate with the WLAN. Some use TKIP, others use CCMP. Since broadcast traffic needs to be understood by all clients, the broadcast encryption type in this scenario is TKIP.

The wireless controller supports WPA/WPA2-TKIP, WPA2-CCMP, WEP 64, WEP 128 and Keyguard encryption options.

Refer to the following to configure an encryption scheme for a wireless controller managed WLAN:

- [WPA/WPA2-TKIP](#)
- [WPA2-CCMP](#)
- [WEP 64](#)
- [WEP 128](#)
- [Keyguard](#)

802.1x EAP, EAP PSK and EAP MAC

[Configuring WLAN Security](#)

The *Extensible Authentication Protocol* (EAP) is the de-facto standard authentication method used to provide secure authenticated access to wireless controller managed WLANs. EAP provides mutual authentication, secured credential exchange, dynamic keying and strong encryption. 802.1X EAP can be deployed with WEP, WPA or WPA2 encryption schemes to further protect user information forwarded over wireless controller managed WLANs.

The EAP process begins when an unauthenticated supplicant (client device) tries to connect with an authenticator (in this case, the authentication server). An Access Point passes EAP packets from the client to an authentication server on the wired side of the access point. All other packet types are blocked until the authentication server (typically, a RADIUS server) verifies the client's identity.

802.1X EAP provides mutual authentication over the WLAN during authentication. The 802.1X EAP process uses credential verification to apply specific policies and restrictions to WLAN users to ensure access is only provided to specific wireless controller resources.

802.1X requires a 802.1X capable RADIUS server to authenticate users and a 802.1X client installed on each device accessing the EAP supported WLAN. An 802.1X client is included with most commercial operating systems, including Microsoft Windows, Linux and Apple OS X.

The RADIUS server authenticating 802.1X EAP users can reside either internally or externally to the RFS4000, RFS6000 or RFS7000 model wireless controller. User account creation and maintenance can be provided centrally using RFMS or individually maintained on each device. If an external RADIUS server is used, EAP authentication requests are forwarded.

When using PSK with EAP, the controller sends a packet requesting a secure link using a pre-shared key. The controller and authenticating device must use the same authenticating algorithm and passcode during authentication. EAP-PSK is useful when transitioning from a PSK network to one that supports EAP. The only encryption types supported with this are TKIP, CCMP and TKIP-CCMP.

To configure EAP on a wireless controller managed WLAN:

1. Select **Configuration > Wireless > Wireless LAN Policy** to display a high-level display of the existing WLANs available to the wireless controller managed network.
2. Select the **Add** button to create an additional WLAN, or select an existing WLAN and **Edit** to modify the security properties of an existing WLAN.
3. Select **Security**.
4. Select **EAP, EAP PSK or EAP-MAC** as the Authentication Type.

Either option enables the radio buttons for various encryption options as an additional measure of security with the WLAN that can be used with EAP.

FIGURE 181 EAP, EAP PSK or EAP MAC Authentication screen

5. Either select an existing **AAA Policy** from the drop-down menu or select the **Create** icon to the right of the AAA Policy parameter to display a screen where new AAA policies can be created. Select the **Edit** icon to modify the configuration of the selected AAA policy.

Authentication, authorization, and accounting (AAA) is a framework for intelligently controlling access to the wireless client managed network, enforcing user authorization policies and auditing and tracking usage. These combined processes are central for securing wireless client resources and wireless network data flows. For information on defining a new AAA policy that can be applied to managed WLAN supporting EAP, EAP PSK or EAP MAC, see *AAA Policy on page 6-289*.

6. Select the **Reauthentication** check box to force EAP supported clients to reauthenticate. Use the spinner control set the number of seconds (between 30 - 86,400) that, once exceeded, forces the EAP supported client to reauthenticate with the controller to use the resources supported by the WLAN.
7. Select **OK** when completed to update the WLAN's EAP configuration. Select **Reset** to revert the screen back to the last saved configuration.

EAP, EAP PSK and EAP MAC Deployment Considerations

802.1x EAP, EAP PSK and EAP MAC

Before defining a 802.1x EAP, EAP PSK or EAP MAC supported configuration on a wireless controller WLAN, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- Brocade Mobility recommends a valid certificate be issued and installed on devices providing 802.1X EAP. The certificate should be issued from an Enterprise or public certificate authority to allow 802.1X clients to validate the identity of the authentication server prior to forwarding credentials.
- If using an external RADIUS server for EAP authentication, Brocade Mobility recommends the round trip delay over the WAN does not exceed 150ms. Excessive delay over a WAN can cause authentication and roaming issues and impact wireless client performance. If experiencing excessive delays, consider using the wireless controllers own RADIUS resources.

MAC Authentication

Configuring WLAN Security

MAC is a device level authentication method used to augment other security schemes when legacy devices are deployed using static WEP.

MAC authentication can be used for device level authentication by permitting WLAN access based on device MAC address. MAC authentication is typically used to augment WLAN security options that do not use authentication (such as static WEP, WPA-PSK and WPA2-PSK) MAC authentication can also be used to assign VLAN memberships, Firewall policies and time and date restrictions.

MAC authentication can only identify devices, not users. MAC authentication only references a client wireless interface card MAC address when authenticating the device, it does not distinguish the device's user credentials. MAC authentication is somewhat poor as a standalone data protection technique, as MAC addresses can be easily spoofed by hackers who can provide a device MAC address to mimic a trusted device within the wireless controller managed network.

With Brocade Mobility RFS4000, RFS6000 and RFS7000 model wireless controllers, MAC authentication is enabled per WLAN profile, augmented with the use of a RADIUS server to authenticate each device. A device's MAC address can be authenticated against the local RADIUS server built into the device or centrally (from a datacenter). For RADIUS server compatibility, the format of the MAC address can be forwarded to the RADIUS server in non-delimited and or delimited formats:

To configure MAC on a wireless controller managed WLAN:

1. Select **Configuration > Wireless > Wireless LAN Policy** to display a high-level display of the existing WLANs available to the wireless controller managed network.
2. Select the **Add** button to create an additional WLAN, or select an existing WLAN and **Edit** to modify the security properties of an existing WLAN.

3. Select **Security**.
4. Select **MAC** as the Authentication Type.

Selecting MAC enables the radio buttons for each encryption option as an additional measure of security for the WLAN.

FIGURE 182 MAC Authentication screen

5. Either select an existing AAA Policy from the drop-down menu or select the **Create** icon to the right of the AAA Policy parameter to display a screen where new AAA policies can be created. A default AAA policy is also available if configuring a WLAN for the first time and there's no existing policies. Select the **Edit** icon to modify the configuration of a selected AAA policy.

Authentication, authorization, and accounting (AAA) is a framework for intelligently controlling access to the wireless client managed network, enforcing user authorization policies and auditing and tracking usage. These combined processes are central for securing wireless client resources and wireless network data flows. For information on defining a new AAA policy that can be applied to managed WLAN supporting MAC, see *AAA Policy on page 6-289*.

6. Select the **Reauthentication** check box to force MAC supported clients to reauthenticate. Use the spinner control set the number of minutes (between 30 - 86,400) that, once exceeded, forces the EAP supported client to reauthenticate with the controller to use the resources supported by the WLAN.
7. Select **OK** when completed to update the WLAN's MAC configuration. Select **Reset** to revert the screen back to the last saved configuration.

MAC Authentication Deployment Considerations

MAC Authentication

Before defining a MAC authentication configuration on a wireless controller WLAN, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- MAC authentication can only be used to identify end-user devices, not the users themselves.
- MAC authentication is somewhat poor as a standalone data protection technique, as MAC addresses can be easily spoofed by hackers who can provision a MAC address on their device to mimic a trusted device.

Kerberos

Configuring WLAN Security

Kerberos (designed and developed by MIT) provides strong authentication for client/server applications using secret-key cryptography. Using Kerberos, a client must prove its identity to a server (and vice versa) across an insecure network connection.

Once a client and server use Kerberos to validate their identity, they encrypt all communications to assure privacy and data integrity. Kerberos can only be used on the access point with Brocade Mobility 802.11b clients.

NOTE

Kerberos makes no provisions for host security. Kerberos assumes that it is running on a trusted host with an untrusted network. If host security is compromised, Kerberos is compromised as well.

Kerberos uses *Network Time Protocol* (NTP) for synchronizing the clocks of its *Key Distribution Center* (KDC) server(s).

To configure Kerberos on a wireless controller managed WLAN:

1. Select **Configuration > Wireless > Wireless LAN Policy** to display a high-level display of the existing WLANs available to the wireless controller managed network.
2. Select the **Add** button to create an additional WLAN, or select an existing WLAN and **Edit** to modify the security properties of an existing WLAN.
3. Select **Security**.
4. Select **Kerberos** as the Authentication Type.

When Kerberos is selected, the **AAA Policy** and **Reauthentication** parameters become disabled, and a **Settings** link displays on the right-hand side of the screen.

5. Select the **Settings** link to define the configuration of the Kerberos supported WLAN.

FIGURE 183 Kerberos Settings screen

KDC

Specify a name that is case-sensitive, for example, BROCADE.COM. The name is the name domain/realm name of the KDC server. A name functions similarly to a DNS domain name. In theory, the name is arbitrary. However, in practice a Kerberos realm is named by upper casing the DNS domain name associated with hosts in the realm. The name must not exceed 127 characters.

KDC Password	Provide the password required to access the KDC server. The password must not exceed 127 characters.
KDC Server Timeout	Specify the time (1 -10 seconds) for the retransmission of Kerberos authentication request packets. If this time is exceeded, the authentication session is retried. The default is 5 seconds.
Primary KDC Server	Specify a numerical (non-DNS) IP address or hostname for the primary <i>Key Distribution Center</i> (KDC). The KDC implements an authentication service and a ticket granting service, whereby an authorized user is granted a ticket encrypted with the user's password. The KDC has a copy of every user's password. Specify the port on which the primary KDC resides. The default port is Port 88.
Secondary KDC Server	Optionally, specify a numerical (non-DNS) IP address or hostname for a secondary remote KDC. Kerberos implementations can use an administration server allowing remote manipulation of the Kerberos database. This administration server usually runs on the KDC. Specify the port on which the secondary KDC resides. The default port is Port 88.

6. Select **OK** when completed to update the WLAN's Kerberos authentication configuration. Select **Reset** to revert the screen back to the last saved configuration.

Kerberos Deployment Considerations

Before defining a Kerberos supported configuration on a wireless controller WLAN, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- Brocade Mobility proprietary authentication techniques such as Kerberos can also be enabled on WLANs supporting legacy KeyGuard supported Symbol clients.
- A Kerberos server's response to an access point contains the client's message and encryption key derived from an EAP-TLS session key. The access point generates a multicast/global authentication key by generating a random number or selecting it from an existing value. On receiving the Kerberos server message, the forwards a success message to wireless client. Consequently, round trip times can be negatively impacted by network congestion.

PSK / None

[Configuring WLAN Security](#)

Open-system authentication can be referred to as no authentication, since no actual authentication takes place. A client requests (and is granted) authentication with no credential exchange.

NOTE

Although None implies no authentication, this option is also used when pre-shared keys are used for encryption (thus the /PSK in the description).

Captive Portal

[Configuring WLAN Security](#)

A *captive portal* is guest access policy for providing guests temporary and restrictive access to the managed wireless network. The primary means of securing such controller guest access is the use of a hotspot. For an overview of the Captive Portal process and information on how to define a captive portal policy that can be applied to a WLAN, see *Configuring Captive Portal Policies on page 10-465*.

To assign a captive portal policy to a managed WLAN:

1. Select **Configuration > Wireless > Wireless LAN Policy** to display a high-level display of the existing WLANs available to the wireless controller managed network.
2. Select the **Add** button to create an additional WLAN or select an existing WLAN and select **Edit** to modify the properties of an existing wireless controller WLAN.
3. Select **Security**.
4. Refer to the **Captive Portal** section within the WLAN Policy security screen




FIGURE 184 WLAN Policy Security screen - Captive Portal Field

5. Select the **Captive Portal Enable** option if authenticated guest access is required with the selected WLAN. This feature is disabled by default.
6. Select the **Captive Portal if Primary Authentication Fails** check box to enable the captive portal policy if the primary authentication is unavailable.
7. Select the **Captive Portal Policy** to use with the WLAN from the drop-down menu. If no relevant policies exist, select the **Create** icon to define a new policy to use with this WLAN or the **Edit** icon to update the configuration of an existing Captive Portal policy. For more information, see [Configuring Captive Portal Policies on page 10-465](#).
8. Select **OK** when completed to update the Captive Portal configuration. Select **Reset** to revert the WLAN Policy Security screen back to the last saved configuration.

WPA/WPA2-TKIP

Configuring WLAN Security

Wi-Fi Protected Access (WPA) is an encryption scheme specified in the IEEE *Wireless Fidelity* (Wi-Fi) standard, 802.11i. WPA provides more sophisticated data encryption than WEP. WPA is designed for corporate networks and small-business environments where more wireless traffic allows quicker discovery of encryption keys by an unauthorized person.

The encryption method is *Temporal Key Integrity Protocol* (TKIP). TKIP addresses WEP's weaknesses with a re-keying mechanism, a per-packet mixing function, a message integrity check, and an extended initialization vector, however TKIP also has vulnerabilities.

Wi-Fi Protected Access 2 (WPA2) is an enhanced version of WPA. WPA2 uses the *Advanced Encryption Standard* (AES) instead of TKIP. AES supports 128-bit, 192-bit and 256-bit keys. WPA/WPA2 also provide strong user authentication based on 802.1x EAP.

To configure WPA/WPA2 encryption on a wireless controller managed WLAN:

1. Select **Configuration > Wireless > Wireless LAN Policy** to display a high-level display of the existing WLANs available to the wireless controller managed network.
2. Select the **Add** button to create an additional WLAN or select an existing WLAN and select **Edit** to modify the properties of an existing wireless controller WLAN.
3. Select **Security**.

- Select the **WPA/WPA2-TKIP** radio button from within the Select Encryption field.

The screen populates with the parameters required to define a WLAN WPA/WPA2-TKIP configuration for the new or existing WLAN.

The screenshot shows the configuration interface for WPA/WPA2-TKIP. At the top, there are four radio buttons: **WPA/WPA2-TKIP** (checked), WPA2-CCMP, WEP 128, and WEP 64. Below these are checkboxes for KeyGuard and an 'Open' button. The 'Key Settings' section includes a 'Pre-Shared Key' field with a dropdown set to 'ASCII'. The 'Key Rotation' section has two rows: 'Unicast Rotation Interval' and 'Broadcast Rotation Interval', both with a checked checkbox and a spinner set to 30 seconds. The 'Fast Roaming' section has three checked options: 'Pairwise Master Key (PMK) Caching', 'Opportunistic Key Caching', and 'Pre-Authentication'. The 'Advanced' section includes 'WPA/WPA2 Handshake Attempts' (3), 'Timeout between Attempts' (four spinners set to 10), 'TKIP Countermeasure Hold Time' (1 Minute), and 'Exclude WPA2 TKIP' (unchecked).

FIGURE 185 WPA/WPA2-TKIP screen

- Define **Key Settings**.

Pre-Shared Key

Enter either an alphanumeric string of 8 to 63 ASCII characters or 64 HEX characters as the primary string both transmitting and receiving authenticators must share. The alphanumeric string allows character spaces. The wireless controller converts the string to a numeric value. This passphrase saves the administrator from entering the 256-bit key each time keys are generated.

- Define **Key Rotation** values.

Unicast messages are addressed to a single device on the network. Broadcast messages are addressed to multiple devices. When using WPA2, a wireless client can use 2 keys, one unicast key, for its own traffic to and from an access point, and one broadcast key, the common key for all the clients in that subnet.

Brocade Mobility recommends rotating the keys so a potential hacker would not have enough data using a single key to attack the deployed encryption scheme.

Unicast Rotation Interval	Define an interval for unicast key transmission in seconds (30 -86,400). Some clients have issues using unicast key rotation, so ensure you know which kind of clients are impacted before using unicast keys. This feature is disabled by default.
Broadcast Rotation Interval	When enabled, the key indices used for encrypting/decrypting broadcast traffic will be alternatively rotated based on the defined interval Define an interval for broadcast key transmission in seconds (30-86,400). Key rotation enhances the broadcast traffic security on the WLAN. This feature is disabled by default.

7. Define the **Fast Roaming** configuration used with the WPA/WPA2-TKIP policy.

Using 802.11i can speed up the roaming process from one AP to another. Instead of doing a complete 802.1x authentication each time a client roams between APs, 802.11i allows a client to re-use previous PMK authentication credentials and perform a four-way handshake. This speeds up the roaming process. In addition to reusing PMKs on previously visited APs, Opportunistic Key Caching allows multiple APs to share PMKs amongst themselves. This allows a client to roam to an AP it has not previously visited and reuse a PMK from another AP to skip 802.1x authentication.

Pairwise Master Key Caching	Select <i>Pairwise Master Key</i> (PMK) caching to store a PMK derived from 802.1x authentication between a client device and its authenticator. When a client roams between devices, the client's credentials no longer need to be completely reauthenticated (a process taking up to 100 milliseconds). With voice sessions, the connection would likely be terminated if not using a PMK. PMK cache entries are stored for a finite amount of time, as configured on the wireless client. When a device initially associates, full 802.1X authentication occurs, and the PMK is cached by the access point and device. If the device roams to a different access point, then roams back, the device already authenticated on the access point, providing faster re-association. This feature is enabled by default.
------------------------------------	--

Opportunistic Key Caching	Opportunistic Key Caching is an extension of PMK caching, allowing a wireless controller to use a PMK derived with a client on one access point with the same client when it roams to another access point. Upon roaming, the client does not have to conduct 802.1x authentication, and can start sending and receiving data sooner. When a device initially associates, full 802.1X authentication occurs and the PMK is cached by the wireless controller and the device. When an authenticated device roams to a different access point managed by the same wireless controller, the device will be already authenticated on the access point providing faster re-association. This feature is enabled by default.
----------------------------------	--

Pre-Authentication	Selecting the Pre-Authentication option enables an associated client to carry out an 802.1x authentication with another wireless controller (or device) before it roams to it. This enables the roaming client to send and receive data sooner by not having to conduct an 802.1x authentication after roaming. With pre authentication, a client can perform an 802.1X authentication with other detected access points while still connected to its current access point. When a device roams to a neighboring access point, the device is already authenticated on the access point providing faster re-association. This feature is disabled by default.
---------------------------	--

- Set the following **Advanced** settings for the WPA/WPA2-TKIP encryption scheme

TKIP Countermeasure Hold Time	The TKIP countermeasure hold-time is the time during which the use of the WLAN is disabled if TKIP countermeasures have been invoked on the WLAN. Use the drop-down menu to define a value in either <i>Hours</i> (0-18), <i>Minutes</i> (0-1,092) or <i>Seconds</i> (0-65,535). The default setting is 60 seconds.
Exclude WPA2 TKIP	Select this option for an Access Point to advertise and enable support for only WPA-TKIP. This option can be used if certain older clients are not compatible with the newer WPA2-TKIP information elements. Enabling this option allows backwards compatibility for clients that support WPA-TKIP and WPA2-TKIP but do not support WPA2-CCMP. Brocade Mobility recommends enabling this feature if WPA-TKIP or WPA2-TKIP supported clients operate in a WLAN populated by WPA2-CCMP enabled clients. This feature is disabled by default.

- Select **OK** when completed to update the WLAN's WPA/WPA2-TKIP encryption configuration. Select **Reset** to revert the screen back to its last saved configuration.

NOTE

WPA-TKIP is not supported on radios configured to exclusively use 802.11n.

WPA-TKIP Deployment Considerations

Before defining a WPA-TKIP supported configuration on a wireless controller WLAN, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- Brocade Mobility recommends TKIP only be enabled for legacy device support when WPA2-CCMP support is not available.
- Though TKIP offers better security than WEP, it can be vulnerable to certain attacks.
- When both TKIP and CCMP are both enabled a mix of clients are allowed to associate with the WLAN. Some use TKIP, others use CCMP. Since broadcast traffic needs to be understood by all clients, the broadcast encryption type in this scenario is TKIP.

WPA2-CCMP

[Configuring WLAN Security](#)

WPA2 is a newer 802.11i standard that provides even stronger wireless security than *Wi-Fi Protected Access* (WPA) and WEP. CCMP is the security standard used by the Advanced Encryption Standard (AES). AES serves the same function TKIP does for WPA-TKIP. CCMP computes a *Message Integrity Check* (MIC) using the proven *Cipher Block Chaining* (CBC) technique. Changing just one bit in a message produces a totally different result.

WPA2/CCMP is based on the concept of a *Robust Security Network* (RSN), which defines a hierarchy of keys with a limited lifetime (similar to TKIP). Like TKIP, the keys the administrator provides are used to derive other keys. Messages are encrypted using a 128-bit secret key and a 128-bit block of data. The end result is an encryption scheme as secure as any the wireless controller provides for its associated clients.

To configure WPA2-CCMP encryption on a wireless controller managed WLAN:

- Select **Configuration > Wireless > Wireless LAN Policy** to display a high-level display of the existing WLANs available to the wireless controller managed network.

2. Select the **Add** button to create an additional WLAN or select an existing WLAN and choose **Edit** to modify the properties of an existing wireless controller WLAN.
3. Select **Security**.
4. Select the **WPA2-CCMP** check box from within the select Select Encryption field.

The screen populates with the parameters required to define a WPA2-CCMP configuration for the new or existing WLAN.

The screenshot displays the WPA2-CCMP configuration interface. At the top, there are four checkboxes: WPA/WPA2-TKIP (unchecked), WPA2-CCMP (checked), WEP 128 (unchecked), KeyGuard (unchecked), WEP 64 (unchecked), and Open (unchecked). Below this is the **Key Settings** section, which includes a text input field for the Pre-Shared Key with a label 'Enter 64 HEX or 8-63 ASCII Characters' and an 'ASCII' dropdown menu. The **Key Rotation** section contains two rows: 'Unicast Rotation Interval' and 'Broadcast Rotation Interval', each with a checkbox (unchecked) and a spinner set to 30, with a range of '(30 to 86,400 seconds)'. The **Fast Roaming** section has three checkboxes: 'Pairwise Master Key(PMK) Caching' (checked), 'Opportunistic Key Caching' (checked), and 'Pre-Authentication' (checked). The **Advanced** section includes 'WPA/WPA2 Handshake Attempts' (spinner set to 3, range '(1 to 5)'), 'Timeout between Attempts' (four spinners, each set to 10, with ranges '(10 - 5000 milliseconds)' for the 3rd and 4th), 'TKIP Countermeasure Hold Time' (input field set to 1, unit 'Minutes', range '(0 to 1,093)'), and 'Exclude WPA2 TKIP' (unchecked).

FIGURE 186 WPA2-CCMP screen

5. Define **Key Settings**.

Pre-Shared Key

Enter either an alphanumeric string of 8 to 63 ASCII characters or 64 HEX characters as the primary string both transmitting and receiving authenticators must share. The alphanumeric string allows character spaces. The wireless controller converts the string to a numeric value. This passphrase saves the administrator from entering the 256-bit key each time keys are generated.

- a. Define **Key Rotation** values.

Unicast messages are addressed to a single device on the network. Broadcast messages are addressed to multiple devices. When using WPA2-CCMP, a wireless client can use 2 keys: one unicast key, for its own traffic to and from an AP, and one broadcast key, the common key for all the clients in that subnet.

Brocade Mobility recommends rotating these keys so a potential hacker would not have enough data using a single key to attack the deployed encryption scheme.

Unicast Rotation Interval	Define an interval for unicast key transmission in seconds (30 -86,400). Some clients have issues using unicast key rotation, so ensure you know which kind of clients are impacted before using unicast keys. This value is disabled by default.
Broadcast Rotation Interval	When enabled, the key indices used for encrypting/decrypting broadcast traffic will be alternatively rotated based on the defined interval Define an interval for broadcast key transmission in seconds (30-86,400). Key rotation enhances the broadcast traffic security on the WLAN. This value is disabled by default.

6. Define the **Fast Roaming** configuration used with the WPA2-CCMP policy.

Using 802.11i can speed up the roaming process from one AP to another. Instead of doing a complete 802.1x authentication each time a client roams between APs, 802.11i allows a client to re-use previous PMK authentication credentials and perform a four-way handshake.

This speeds up the roaming process. In addition to reusing PMKs on previously visited APs, Opportunistic Key Caching allows multiple APs to share PMKs amongst themselves. This allows a client to roam to an AP it has not previously visited and reuse a PMK from another AP to skip 802.1x authentication.

Pairwise Master Key (PMK) Caching	Select Pairwise Master Key (PMK) Caching to store a PMK derived from 802.1x authentication between a client device and its authenticator. When a client roams between devices, the client's credentials no longer need to be completely reauthenticated (a process taking up to 100 milliseconds). With voice sessions, the connection would likely be terminated if not using a PMK. PMK cache entries are stored for a finite amount of time, as configured on the wireless client. When a device initially associates, full 802.1X authentication occurs, and the PMK is cached by the access point and device. If the device roams to a different access point, then roams back, the device already authenticated on the access point, providing faster re-association. This feature is enabled by default.
--	--

Opportunistic Key Caching	Opportunistic Key Caching is an extension of PMK caching, allowing a wireless controller to use a PMK derived with a client on one access point with the same client when it roams to another access point. Upon roaming, the client does not have to conduct 802.1x authentication, and can start sending and receiving data sooner. When a device initially associates, full 802.1X authentication occurs and the PMK is cached by the wireless controller and the device. When an authenticated device roams to a different access point managed by the same wireless controller, the device will be already authenticated on the access point providing faster re-association. This feature is enabled by default.
----------------------------------	--

Pre-Authentication	Selecting the Pre-Authentication option enables an associated client to carry out an 802.1x authentication with another wireless controller (or device) before it roams to it. This enables the roaming client to send and receive data sooner by not having to conduct an 802.1x authentication after roaming. With pre authentication, a wireless client can perform an 802.1X authentication with other detected access points while still connected to its current access point. When a device roams to a neighboring AP, the device is already authenticated on the access point providing faster re-association. This feature is enabled by default.
---------------------------	--

- Set the following **Advanced** for the WPA2-CCMP encryption scheme.

TKIP Countermeasure Hold Time The TKIP countermeasure hold-time is the time during which the use of the WLAN is disabled if TKIP countermeasures have been invoked on the WLAN. Use the drop-down menu to define a value in either *Hours* (0-18), *Minutes* (0-1,092) or *Seconds* (0-65,535). The default setting is 60 seconds.

Exclude WPA2-TKIP Select this option for an Access Point to advertise and enable support for only WPA-TKIP. Select this option if certain older clients are not compatible with the newer WPA2-TKIP information elements. Enabling this option allows backwards compatibility for clients that support WPA-TKIP and WPA2-TKIP but do not support WPA2-CCMP. Brocade Mobility recommends enabling this feature if WPA-TKIP or WPA2-TKIP supported clients operate in a WLAN populated by WPA2-CCMP enabled clients. This feature is disabled by default.

- Select **OK** when completed to update the WLAN's WPA2-CCMP encryption configuration. Select **Reset** to revert back to its last saved configuration.

WPA2-CCMP Deployment Considerations

WPA2-CCMP

Before defining a WPA2-CCMP supported configuration on a wireless controller WLAN, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- Brocade Mobility recommends WPA2-CCMP be configured for all new (non visitor) WLANs requiring encryption, as it's supported by the majority of the hardware and client vendors using Brocade Mobility wireless networking equipment.
- WPA2-CCMP supersedes WPA-TKIP and implements all the mandatory elements of the 802.11i standard. WPA2-CCMP introduces a new AES-based algorithm called CCMP which replaces TKIP and WEP and is considered significantly more secure.

WEP 64

Configuring WLAN Security

Wired Equivalent Privacy (WEP) is a security protocol specified in the IEEE *Wireless Fidelity* (Wi-Fi) standard. WEP is designed to provide a WLAN with a level of security and privacy comparable to that of a wired LAN.

WEP can be used with open, shared, MAC and 802.1 X EAP authentications. WEP is optimal for WLANs supporting legacy deployments when also used with 802.1X EAP authentication to provide user and device authentication and dynamic WEP key derivation and periodic key rotation. 802.1X provides authentication for devices and also reduces the risk of a single WEP key being deciphered. If 802.1X support is not available on the legacy device, MAC authentication should be enabled to provide device level authentication.

WEP 64 uses a 40 bit key concatenated with a 24-bit initialization vector (IV) to form the RC4 traffic key. WEP 64 is a less robust encryption scheme than WEP 128 (containing a shorter WEP algorithm for a hacker to potentially duplicate), but networks that require more security are at risk from a WEP flaw. WEP is only recommended if there are client devices that are incapable of using higher forms of security. The existing 802.11 standard alone offers administrators no effective method to update keys.

To configure WEP 64 encryption on a managed WLAN:

- Select **Configuration > Wireless > Wireless LAN Policy** to display a high-level display of the existing WLANs available to the wireless controller managed network.

2. Select the **Add** button to create an additional WLAN or select **Edit** to modify the properties of an existing wireless controller WLAN.
3. Select **Security**.
4. Select the **WEP 64** check box from within the Select Encryption field.

The screen populates with the parameters required to define a WEP 64 configuration for the WLAN.

The screenshot shows the 'Select Encryption' configuration interface. At the top, there are four checkboxes: 'WPA/WPA2-TKIP', 'WPA2-CCMP', 'WEP 128', and 'KeyGuard'. The 'WEP 64' checkbox is checked. To the right of these is an 'Open' checkbox. Below this is a text input field labeled 'Enter 4 to 32 Characters' and a 'Generate Keys' button. Underneath, there are four rows for 'Key 1' through 'Key 4'. Each row has a text input field labeled 'Enter 10 HEX or 5 ASCII Characters', a 'HEX' dropdown menu, and a radio button under the heading 'Transmit Key'. At the bottom of the key configuration area is a 'Restore Default WEP Keys' button.

FIGURE 187 WEP 64 screen

5. Configure the following WEP 64 settings:

Generate Keys

Specify a 4 to 32 character Pass Key and click the **Generate** button. The pass key can be any alphanumeric string. The wireless controller, other proprietary routers, and Brocade Mobility clients use the algorithm to convert an ASCII string to the same hexadecimal number. Clients without Brocade Mobility adapters need to use WEP keys manually configured as hexadecimal numbers.

Keys 1-4

Use the Key #1-4 fields to specify key numbers. For WEP 64 (40-bit key), the keys are 10 hexadecimal characters in length. Select one of these keys for default activation by clicking its radio button.

Restore Default WEP Keys

If you feel it necessary to restore the WEP algorithm back to its default settings, click the **Restore Default WEP Keys** button.

Default WEP 64 keys are as follows:

- Key 1 1011121314
- Key 2 2021222324
- Key 3 3031323334
- Key 4 4041424344

6. Select **OK** when completed to update the WLAN's WEP 64 encryption configuration. Select **Reset** to revert the screen back to its last saved configuration.

WEP 64 Deployment Considerations

Before defining a WEP 64 supported configuration on a wireless controller WLAN, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- Brocade Mobility recommends additional layers of security (beyond WEP) be enabled to minimize the likelihood of data loss and security breaches. WEP enabled WLANs should be mapped to an isolated VLAN with Firewall policies restricting access to hosts and suspicious network applications.
- WEP enabled WLANs should only be permitted access to resources required by legacy devices.
- If WEP support is needed for WLAN legacy device support, 802.1X EAP authentication should be also configured in order for the WLAN to provide authentication and dynamic key derivation and rotation.

WEP 128

Configuring WLAN Security

Wired Equivalent Privacy (WEP) is a security protocol specified in the IEEE *Wireless Fidelity (Wi-Fi)* standard. WEP is designed to provide a WLAN with a level of security and privacy comparable to that of a wired LAN.

WEP can be used with open, shared, MAC and 802.1 X EAP authentications. WEP is optimal for WLANs supporting legacy deployments when also used with 802.1X EAP authentication to provide user and device authentication and dynamic WEP key derivation and periodic key rotation. 802.1X provides authentication for devices and also reduces the risk of a single WEP key being deciphered. If 802.1X support is not available on the legacy device, MAC authentication should be enabled to provide device level authentication.

WEP 128 uses a 104 bit key which is concatenated with a 24-bit initialization vector (IV) to form the RC4 traffic key. WEP may be all a small-business user needs for the simple encryption of wireless data. However, networks that require more security are at risk from a WEP flaw. WEP is only recommended if there are client devices that are incapable of using higher forms of security. The existing 802.11 standard alone offers administrators no effective method to update keys.

WEP 128 provides a more robust encryption algorithm than WEP 64 by requiring a longer key length and pass key. Thus, making it harder to hack through the replication of WEP keys.

To configure WEP 128 encryption on a wireless controller managed WLAN:

1. Select **Configuration > Wireless > Wireless LAN Policy** to display a high-level display of the existing WLANs available to the wireless controller managed network.
2. Select the **Add** button to create an additional WLAN or select **Edit** to modify the properties of an existing wireless controller WLAN.
3. Select **Security**.
4. Select the **WEP 128** check box from within the Select Encryption field.

The screen populates with the parameters required to define a WEP 128 configuration for the WLAN.

Select Encryption

WPA/WPA2-TKIP WEP 128 WEP 64 Open
 WPA2-CCMP KeyGuard

Enter 4 to 32 Characters

Generate Keys

Enter 26 HEX or 13 ASCII Characters

Key 1 HEX

Key 2 HEX

Key 3 HEX

Key 4 HEX

Transmit Key

FIGURE 188 WEP 128 screen

- Configure the following WEP 128 settings:

Generate Keys

Specify a 4 to 32 character Pass Key and click the **Generate** button. The pass key can be any alphanumeric string. The wireless controller, other proprietary routers, and Brocade Mobility clients use the algorithm to convert an ASCII string to the same hexadecimal number. Clients without Brocade Mobility adapters need to use WEP keys manually configured as hexadecimal numbers.

Keys 1-4

Use the Key #1-4 areas to specify key numbers. For WEP 128 (104-bit key), the keys are 26 hexadecimal characters in length. Select one of these keys for default activation by clicking its radio button.

Restore Default WEP Keys

If you feel it necessary to restore the WEP algorithm back to its default settings, click the **Restore Default WEP Keys** button.

Default WEP 128 keys are as follows:

- Key 1 101112131415161718191A1B1C
- Key 2 202122232425262728292A2B2C
- Key 3 303132333435363738393A3B3C
- Key 4 404142434445464748494A4B4C

- Select **OK** when completed to update the WLAN's WEP 128 encryption configuration. Select **Reset** to revert the screen back to its last saved configuration.

WEP 128 Deployment Considerations

WEP 128

Before defining a WEP 128 supported configuration on a wireless controller WLAN, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- Brocade Mobility recommends additional layers of security (beyond WEP) be enabled to minimize the likelihood of data loss and security breaches. WEP enabled WLANs should be mapped to an isolated VLAN with Firewall policies restricting access to hosts and suspicious network applications.
- WEP enabled WLANs should only be permitted access to resources required by legacy devices.
- If WEP support is needed for WLAN legacy device support, 802.1X EAP authentication should be also configured in order for the WLAN to provide authentication and dynamic key derivation and rotation.

Keyguard

Configuring WLAN Security

Keyguard is a form of WEP, and could be all a small business needs for the simple encryption of wireless data.

KeyGuard is a proprietary encryption method developed by Brocade Mobility. KeyGuard is Brocade Mobility's enhancement to WEP encryption, and was developed before the finalization of WPA-TKIP. The Keyguard encryption implementation is based on the IEEE Wi-Fi standard, 802.11i.

To configure Keyguard encryption on a wireless controller managed WLAN:

1. Select **Configuration > Wireless > Wireless LAN Policy** to display a high-level display of the existing WLANs available to the wireless controller managed network.
2. Select the **Add** button to create an additional WLAN or select **Edit** to modify the properties of an existing wireless controller WLAN.
3. Select **Security**.
4. Select the **Keyguard** check box from within the Select Encryption field.

The screen populates with the parameters required to define a KeyGuard configuration for the WLAN.

FIGURE 189 WLAN KeyGuard Configuration screen

5. Configure the following Keyguard settings:

Generate Keys

Specify a 4 to 32 character Pass Key and click the **Generate** button. The pass key can be any alphanumeric string. The wireless controller or other proprietary routers. Brocade Mobility clients use the algorithm to convert an ASCII string to the same hexadecimal number. Clients without Brocade Mobility adapters need to use keys manually configured as hexadecimal numbers.

Keys 1-4

Use the Key #1-4 areas to specify key numbers. For Keyguard (104-bit key), the keys are 26 hexadecimal characters in length. Select one of these keys for default activation by clicking its radio button.

Restore Default WEP Keys

If you feel it necessary to restore the Keyguard algorithm back to its default settings, click the **Restore Default WEP Keys** button. This may be the case if the latest defined algorithm has been compromised and no longer provides its former measure of data security.

Default WEP Keyguard keys are as follows:

- Key 1 101112131415161718191A1B1C
- Key 2 202122232425262728292A2B2C
- Key 3 303132333435363738393A3B3C
- Key 4 404142434445464748494A4B4C

6. Select **OK** when completed to update the WLAN's Keyguard encryption configuration. Select **Reset** to revert the screen back to its last saved configuration.

KeyGuard Deployment Considerations

Keyguard

Before defining a Keyguard configuration on a WLAN, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- Brocade Mobility proprietary authentication techniques, such as Kerberos, can also be enabled on WLANs supporting other Brocade Mobility proprietary techniques, such as KeyGuard.
- A wireless controller WLAN using KeyGuard to support legacy Brocade Mobility devices should also use largely limited to the support of just those legacy clients using KeyGuard.

Configuring WLAN Firewall Support

Wireless LAN Policy

A Firewall is a mechanism enforcing access control, and is considered a first line of defense in protecting proprietary information within the Brocade Mobility wireless controller managed network. The means by which this is accomplished varies, but in principle, a Firewall can be thought of as mechanisms both blocking and permitting data traffic within the wireless controller managed network. For an overview of Firewalls, see [Wireless Firewall on page 9-419](#).

managed WLANs use Firewalls like *Access Control Lists (ACLs)* to filter/mark packets based on the WLAN from which they arrive, as opposed to filtering packets on Layer 2 ports. An ACL contains an ordered list of *Access Control Entries (ACEs)*. Each ACE specifies an action and a set of conditions (rules) a packet must satisfy to match the ACE. The order of conditions in the list is critical because the wireless controller stops testing conditions after the first match.

IP based Firewall rules are specific to source and destination IP addresses and the unique rules and precedence orders assigned. Both IP and non-IP traffic on the same Layer 2 interface can be filtered by applying both an IP ACL and a MAC.

Additionally, the controller allows administrators to filter Layer 2 traffic on a physical Layer 2 interface using MAC addresses. A MAC Firewall rule uses source and destination MAC addresses for matching operations, where the result is a typical allow, deny or mark designation to WLAN controller packet traffic.

Keep in mind IP and non-IP traffic on the same Layer 2 interface can be filtered by applying both an IP ACL and a MAC ACL to the interface.

To review access policies, create a new policy or edit the properties of an existing policy:

1. Select **Configuration > Wireless LANs > Wireless LAN Policy** to display a high-level display of the existing WLANs available to the wireless controller managed network.
2. Select the **Add** button to create a new WLAN or **Edit** to modify the properties of an existing wireless controller WLAN.
3. Select **Firewall** from the Wireless LAN Policy options.

IP Firewall Rules

Inbound IP Firewall Rules: IPACLRule1

Outbound IP Firewall Rules: IPACLRule1

MAC Firewall Rules

Inbound MAC Firewall Rules: [Empty]

Outbound MAC Firewall Rules: [Empty]

Association ACL

Association ACL: [Empty]

Trust Parameters

ARP Trust:

Validate ARP Header Mismatch:

DHCP Trust:

Wireless Client Deny

Wireless Client Denied Traffic Threshold: 1 (1 to 1,000,000 packets per second)

Action: None

Blacklist Duration: 0 (0 to 86,400 seconds)

Advanced

Firewall Session Hold Time: 30 Seconds (1 to 300)

OK Reset Exit

FIGURE 190 WLAN Policy Firewall screen

The screen displays editable fields for IP Firewall Rules, MAC Firewall Rules, Trust Parameters and Client Deny Limits.

Select an existing **Inbound IP Firewall Rule** and **Outbound IP Firewall Rule** using the drop-down menu. If no rules exist, select the **Create** icon to display a screen where Firewall rules can be created. Select the **Edit** icon to modify the configuration of a selected Firewall policy configuration.

4. If creating a new rule, provide a name up to 64 characters in length.
5. Select the **+ Add Row** button.
6. Select the added row to expand it into configurable parameters.

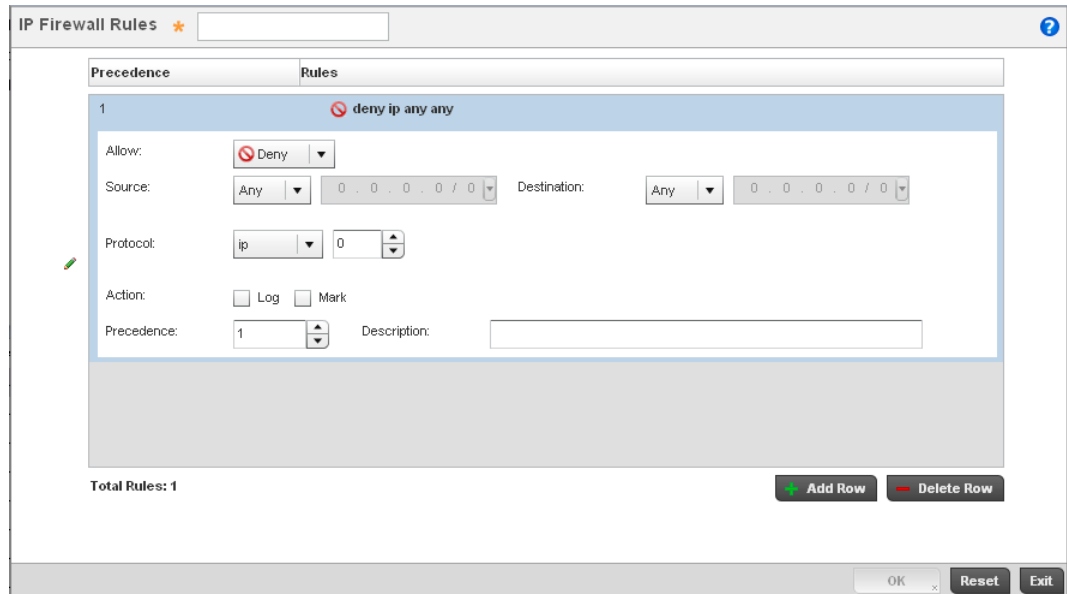


FIGURE 191 IP Firewall Rules screen

7. Define the following parameters for either the inbound or outbound IP Firewall Rules:

Allow	<p>Every IP Firewall rule is made up of matching criteria rules. The action defines what to do with the packet if it matches the specified criteria. The following actions are supported:</p> <p><i>Deny</i>—Instructs the Firewall to prohibit a packet from reaching its destination.</p> <p><i>Permit</i>—Instructs the Firewall to allow a packet to proceed to its destination.</p>
Source and Destination	<p>Enter both Source and Destination IP addresses. The wireless controller uses the source IP address, destination IP address and IP protocol type as basic matching criteria. The wireless controller's access policy filter can also include other parameters specific to a protocol type (like source and destination port for TCP/UDP protocol. Provide a subnet mask if needed.</p>
Protocol	<p>Select the protocol used with the IP access policy from the drop-down menu. IP is selected by default. Selecting <i>ICMP</i> displays an additional set of ICMP specific Options for ICMP Type and code. Selecting either <i>TCP</i> or <i>UDP</i> displays an additional set of specific TCP/UDP source and destinations port options.</p>
Action	<p>The following actions are supported:</p> <p><i>Log</i>—Creates a log entry that a Firewall rule has allowed a packet to either be denied or permitted.</p> <p><i>Mark</i>—Modifies certain fields inside the packet, then permits them. Therefore, mark is an action with an implicit permit.</p> <p><i>Mark, Log</i> — Conducts both mark and log functions.</p>
Precedence	<p>Use the spinner control to specify a precedence for this IP policy between 1-1500. Rules with lower precedence are always applied first to packets.</p>
Description	<p>Provide a description up to characters long for rule to help differentiate it from others with similar configurations.</p>

8. Select existing inbound and outbound **MAC Firewall Rules** using the drop-down menu. If no rules exist, select **Create** to display a screen where Firewall rules can be created.

9. Select the **+ Add Row** button.

10. Select the added row to expand it into configurable parameters.

The screenshot displays the 'MAC Firewall Rules' configuration interface. At the top, it shows 'MAC Firewall Rules test1'. Below this, there are two tabs: 'Precedence' and 'Rules'. The 'Rules' tab is active, showing a table with one rule. The rule is named 'deny any any' and is expanded to show its configuration parameters:

- Allow:** Deny (selected)
- VLAN ID:** 1
- Match 802.1P:** 0
- Source MAC:** Any
- Destination MAC:** Any
- Action:** Log and Mark (checked)
- Ethertype:** other
- Precedence:** 1
- Description:** (empty text box)

At the bottom of the screen, there are buttons for 'Add Row' and 'Delete Row', and a 'Total Rules: 1' indicator. The bottom right corner has 'OK' and 'Reset' buttons.

FIGURE 192 MAC Firewall Rules screen

11. Define the following parameters for either the inbound or outbound MAC Firewall Rules:

- Allow** Every IP Firewall rule is made up of matching criteria rules. The action defines what to do with the packet if it matches the specified criteria. The following actions are supported:
- Deny*—Instructs the Firewall to not to allow a packet to proceed to its destination.
 - Permit*—Instructs the Firewall to allows a packet to proceed to its destination.
- VLAN ID** Enter a VLAN ID representative of the shared SSID each user employs to interoperate within the managed network (once authenticated by the local RADIUS server). The VLAN ID can be between 1 - 4094.
- Match 802.1P** Configures IP DSCP to 802.1p priority mapping for untagged frames. Use the spinner control to define a setting between 0-7.
- Source and Destination MAC** Enter both **Source** and **Destination** MAC addresses. The wireless controller uses the source IP address, destination MAC address as basic matching criteria. Provide a subnet mask if using a mask.
- Action** The following actions are supported:
- Log*—Creates a log entry that a Firewall rule has allowed a packet to either be denied or permitted.
 - Mark*—Modifies certain fields inside the packet and then permits them. Therefore, mark is an action with an implicit permit.
 - Mark, Log* — Conducts both mark and log functions.

Ethertype	Use the drop-down menu to specify an EtherType of either ipv6, arp, wisp, monitor 8021q. An EtherType is a two-octet field within an Ethernet frame. It is used to indicate which protocol is encapsulated in the payload of an Ethernet frame.
Precedence	Use the spinner control to specify a precedence for this MAC Firewall rule between 1-1500. Access policies with lower precedence are always applied first to packets.
Description	Provide a description (up to 64 characters) for the rule to help differentiate the it from others with similar configurations.

12. If creating an new **Association ACL**, provide a name specific to its function. Avoid naming it after a WLAN it may support. The name cannot exceed 32 characters.

13. Save the changes to the new MAC rule or reset to the last saved configuration as needed.

14. Set the following **Trust Parameters**:

ARP Trust	Select the check box to enable ARP Trust on this WLAN. ARP packets received on this managed WLAN are considered trusted and information from these packets is used to identify rogue devices within the managed network. This setting is disabled by default.
Validate ARP Header Mismatch	Select this option to verify the mismatch for source MAC in the ARP and Ethernet headers. By default, mismatch verification is enabled.
DHCP Trust	Select the check box to enable DHCP trust on this WLAN. This setting is disabled by default.

15. Set the following **Wireless Client Deny** configuration:

Wireless Client Denied Traffic Threshold	If enabled, any associated client which exceeds the thresholds configured for storm traffic is either deauthenticated or blacklisted depending on the selected Action. The threshold range is 1-1000000 packets per second. This feature is disabled by default.
Action	If enabling a wireless client threshold, use the drop-down menu to determine whether clients are deauthenticated when the threshold is exceeded or blacklisted from controller connectivity for a user defined interval. Selecting None applies no consequence to an exceeded threshold.
Blacklist Duration	Select the checkbox and define a setting between 0 - 86,400 seconds. Once the blacklist duration has been exceeded, offending clients can reauthenticate with the controller.

16. Set a **Firewall Session Hold Time** in either *Seconds* (1 - 300) or *Minutes* (1 - 5). This is the hold time for caching user credentials and Firewall state information when a client roams. The default setting is 30 seconds.

17. Select **OK** when completed to update this WLAN's Firewall settings. Select **Reset** to revert the screen back to its last saved configuration.

WLAN Firewall Deployment Considerations

Before defining an access control configuration on a wireless controller WLAN, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- IP and non-IP traffic on the same Layer 2 interface can be filtered by applying both an IP ACL and a MAC ACL to the interface.

Configuring Client Settings

Wireless LAN Policy

Each managed WLAN can maintain its own client setting configuration. These include wireless client inactivity timeouts and broadcast configurations.

1. Select **Configuration > Wireless > Wireless LAN Policy** to display a high-level display of the existing WLANs available to the wireless controller managed network.
2. Select the **Add** button to create an additional WLAN, or select an existing WLAN and **Edit** to modify the properties of an existing WLAN.
3. Select the **Client Settings** tab.

Client Settings

- Enable Client-to-Client Communication
- Wireless Client Power (0 to 20 dBm)
- Wireless Client Idle Time Minutes (1 to 1,440)
- Max Firewall Sessions per Client (10 to 10,000)
- Enforce Client Load Balancing
- Enforce DHCP Client Only
- Proxy ARP Mode
- Enforce DHCP Offer Validation

Motorola Solutions Client Extensions

- Move Operations
- Smart Scan
- Symbol Information Element
- WMM Load Information Element

Timeout Settings

- Credential Cache Timeout Days (0 to 1)
- VLAN Cache Timeout Hours (0 to 24)

OK Reset Exit

FIGURE 193 WLAN Policy Client Settings screen

4. Define the following **Client Settings** for the WLAN:

Disallow Client-to-Client Communication	Select this option to disallow client to client communication within this WLAN. The default is enabled, meaning clients are allowed to exchange packets with other clients. It does not necessarily prevent clients on other WLANs from sending packets to this WLAN, but as long as this setting also disabled on that WLAN, the clients are not permitted to interoperate.
Wireless Client Power	Use this parameter to set the maximum transmit power (between 0 - 20 dBm) communicated to wireless clients for transmission within the managed network. The default value is 20 dBm.
Wireless Client Idle Time	Set the maximum amount of time wireless clients are allowed to be idle within this WLAN. Set the idle time in either <i>Seconds</i> (60 - 86,400), <i>Minutes</i> (1 - 1,440), <i>Hours</i> (0 - 24) or <i>Days</i> (0 - 1). When this setting is exceeded, the client is no longer able to access controller resources and must re-authenticate. The default value is 1,800 seconds.
Max Firewall Sessions per Client	Select this option to set the maximum amount of sessions (between 10 - 10,000) clients within the managed network over the Firewall. When enabled, this parameter limits the number of simultaneous sessions allowed by the Firewall per wireless client. This feature is disabled by default.
Enforce Client Load Balancing	Select the checkbox to distribute all the managed clients evenly amongst the Access Point radios associated with the controller. This feature is disabled by default.
Enforce DHCP Client Only	Select the checkbox to enforce that the firewall only allows packets from clients if they used DHCP to obtain an IP address, disallowing static IP addresses. This feature is disabled by default.
Proxy ARP Mode	Use the drop-down menu to define the proxy ARP mode as either Strict or Dynamic. Proxy ARP is the technique used by the AP to answer ARP requests intended for another system. By faking its identity, the AP accepts responsibility for routing packets to the actual destination. Dynamic is the default value.
Enforce DHCP-Offer Validation	Select the checkbox to enforce DHCP offer validation. The default setting is disabled.

5. Define the following **Brocade Mobility Client Extensions** for the WLAN:

Move Operation	Select the checkbox to enable the use of <i>Fast Roaming</i> (HFSR) for clients on this WLAN. This feature applies only to certain Brocade Mobility client devices. This feature is disabled by default.
Smart Scan	Enable a smart scan to refine a clients channel scans to just a few channels as opposed to all available channels. This feature is disabled by default.
Symbol Information Element	Select the checkbox to support the Symbol Information Element with legacy Symbol Technology clients. The default setting is enabled.
WMM Load Information Element	Select the checkbox to support a WMM Load Information Element in radio transmissions with legacy Brocade Mobility clients. The default setting is disabled.

6. Define the following **Timeout Settings** for the WLAN:

Credential Cache Timeout	Set a timeout period for the credential cache in <i>Days</i> , <i>Hours</i> , <i>Minutes</i> or <i>Seconds</i> .
VLAN Cache Timeout	Set a timeout period for the VLAN cache in <i>Days</i> , <i>Hours</i> , <i>Minutes</i> or <i>Seconds</i> .

7. Select **OK** when completed to update the WLAN's client setting configuration. Select **Reset** to revert the screen back to the last saved configuration.

WLAN Client Setting Deployment Considerations

Configuring Client Settings

Before defining a WLAN's client settings, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- Clients on the same WLAN associated to an AAP can communicate locally at the AP Level without going through the controller. If this is undesirable, an Access Point's **Client-to-Client Communication** option should be disabled.
- When the wireless client idle time setting is exceeded, the client is no longer able to access controller WLAN resources and must re-authenticate. The default value is 1,800 seconds.

Configuring WLAN Accounting Settings

Accounting is the method of collecting and sending security server information for billing, auditing, and reporting user data; such as start and stop times, executed commands (such as PPP), number of packets and number of bytes. Accounting enables wireless network administrators to track the services users are accessing and the network resources they are consuming. When accounting is enabled, the network access server reports and logs user activity to a RADIUS security server in the form of accounting records. Each accounting record is comprised of AV pairs and is stored on the controller's access control server. The data can be analyzed for network management, client billing, and/or auditing. Accounting methods must be defined through AAA.

Accounting can be enabled and applied to managed WLANs, to uniquely log accounting events specific to the controller WLAN. Accounting logs contain information about the use of remote access services by users. This information is of great assistance in partitioning local versus remote users and how to best accommodate each. Remote user information can be archived to a location outside of the controller for periodic network and user permission administration.

To configure WLAN accounting settings:

1. Select **Configuration > Wireless LANs > Wireless LAN Policy** to display a high-level display of the existing WLANs available to the wireless controller managed network.
2. Select the **Add** button to create an additional WLAN or select **Edit** to modify the properties of an existing wireless controller WLAN.
3. Select the **Add** button to create an additional WLAN or select **Edit** to modify the properties of an existing wireless controller WLAN.
4. Select **Accounting**.

FIGURE 194 WLAN Policy Accounting screen

- Set the following **System Log Accounting** information:

Enable Syslog Accounting

Use this option for the controller or Access Point to generate accounting records in standard syslog format (RFC 3164). The feature is disabled by default.

Syslog Host

Specify the IP address or hostname of the external syslog host where accounting records are routed.

Syslog Port

Use the spinner control to set the destination UDP port number of the external syslog host where the accounting records are routed.

- Select the **Enable RADIUS Accounting** check box to use an external RADIUS resource for AAA accounting. When the check box is selected, a AAA Policy field displays. Either use the default AAA policy with the WLAN, or select **Create** to define a new AAA configuration that can be applied to the WLAN. This setting is disabled by default.
- Select **OK** when completed to update this WLAN's accounting settings. Select **Reset** to revert the screen to its last saved configuration.

Accounting Deployment Considerations

Before defining a AAA configuration on a controller WLAN, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- When using RADIUS authentication, Brocade Mobility recommends the WAN port round trip delay not exceed 150ms. Excessive delay over a WAN can cause authentication and roaming issues. When excessive delays exist, a distributed RADIUS service should be used.
- Brocade Mobility recommends authorization policies be implemented when users need to be restricted to specific WLANs, or time and date restrictions need to be applied.

- Authorization policies can also apply bandwidth restrictions and assign Firewall policies to users and devices.

Configuring Client Load Balancing Settings

Wireless LAN Policy

To configure advanced settings on a wireless controller managed WLAN:

1. Select **Configuration > Wireless LANs > Wireless LAN Policy** to display a high-level display of the existing WLANs available to the wireless controller managed network.
2. Select the **Add** button to create an additional WLAN or select **Edit** to modify the properties of an existing wireless controller WLAN.
3. Select **Client Load Balancing**.

The screenshot shows the 'Client Load Balancing' configuration window. It is organized into three main sections:

- Load Balancing Settings:**
 - Enforce Client Load Balancing:** Checked (checkbox).
 - Band Discovery Interval:** 24 (input field), Seconds (dropdown), (0 to 10,000) (range).
 - Capability Ageout Time:** 24 (input field), Seconds (dropdown), (0 to 10,000) (range).
- Load Balancing Settings (2.4GHz):**
 - Allow Single Band Client:** Checked (checkbox).
 - Max Probe Requests:** 48 (input field), (0 to 10,000) (range).
 - Probe Request Interval:** 24 (input field), Seconds (dropdown), (0 to 10,000) (range).
- Load Balancing Settings (5GHz):**
 - Allow Single Band Client:** Checked (checkbox).
 - Max Probe Requests:** 24 (input field), (0 to 10,000) (range).
 - Probe Request Interval:** 24 (input field), Seconds (dropdown), (0 to 10,000) (range).

At the bottom right, there are three buttons: 'OK', 'Reset', and 'Exit'.

FIGURE 195 WLAN Policy Client Load Balancing screen

4. Refer to the **Load Balancing Settings** section to configure load balancing for the WLAN.

Enforce Client Load Balancing Select this option to enable client load balancing for the selected WLAN.

5. Refer to the **Load Balancing Settings (2.4GHz)** section to configure load balancing for the 2.4 GHz WLAN.

Allowed Single Band Clients	Select this option to enable association for single band clients on the 2.4GHz frequency, even if load balancing is available.
Max Probe Requests	Enter a value between 0 and 10,000 for the maximum number of probe requests for clients using the 2.4GHz frequency. The default value is 60.
Probe Request Interval	Enter a value in seconds between 0 and 10,000 to configure the time interval for client probe requests beyond which it is allowed to associate for clients on the 2.4GHz network.
Band Discovery Interval	Enter a value in seconds between 0 and 10,000 to configure the time interval to discover client's band capability before associating it

6. Refer to the **Load Balancing Settings (5GHz)** section to configure load balancing for the 5 GHz WLAN.

Allowed Single Band Clients	Select this option to enable association for single band clients on the 5GHz frequency, even if load balancing is available.
Max Probe Requests	Enter a value between 0 and 10,000 for the maximum number of probe requests for clients using the 5GHz frequency. The default value is 60.
Probe Request Interval	Enter a value in seconds between 0 and 10,000 to configure the time interval for client probe requests beyond which it is allowed to associate for clients on the 5GHz network.
Band Discovery Interval	Enter a value in seconds between 0 and 10,000 to configure the time interval to discover client's band capability before associating it

7. Select **OK** when completed to update this WLAN's advanced settings. Select **Reset** to revert the screen back to its last saved configuration.

Configuring Advanced WLAN Settings

Wireless LAN Policy

To configure advanced settings on a wireless controller managed WLAN:

1. Select **Configuration > Wireless LANs > Wireless LAN Policy** to display a high-level display of the existing WLANs available to the wireless controller managed network.
2. Select the **Add** button to create an additional WLAN or select **Edit** to modify the properties of an existing wireless controller WLAN.
3. Select **Advanced**.

The screenshot displays the 'WLAN Policy Advanced' configuration screen. It is organized into three main sections:

- Protected Management Frames:**
 - Mode:** Three radio buttons are present: 'Disabled' (selected), 'Optional', and 'Mandatory'.
 - SA Query Attempts:** A spinner control is set to '3', with a range of '(1 to 15)'.
 - SA Query Retry Timeout:** A spinner control is set to '1000', with a range of '(100 to 6,000 milliseconds)'.
- Advanced RADIUS Configuration:**
 - NAS Identifier:** An empty text input field.
 - NAS Port:** An empty text input field.
 - RADIUS Dynamic Authorization:** A checkbox that is currently unchecked.
- Radio Rates:**
 - Rates for 2.4 GHz WLAN:** A dropdown menu set to 'Custom:' with a 'Select' button.
 - Rates for 5 GHz WLAN:** A dropdown menu set to 'Custom:' with a 'Select' button.

At the bottom right of the screen, there are three buttons: 'OK', 'Reset', and 'Exit'.

FIGURE 196 WLAN Policy Advanced screen

4. Refer to the **Protected Management Frames** field to set a frame protection mode and security association for the WLAN's advanced configuration.

During a *security association (SA)* negotiation, the wireless controller and recipient gateways agree to use a particular transform set to protect data flow. A transform set is a combination of security protocols and algorithms. During an IPSec security association negotiation, peers agree to use a particular transform set for protecting the wireless controller managed data flow.

Mode	Select a radio button option to determine whether management frames are continually or optionally protected. Disabled is the default setting.
SA Query Attempts	Use the spinner control to set the number of security association query attempts between 1-15. The default value is 3.
SA Query Retry Timeout	The timeout value is the configurable interval used to timeout association requests that exceed the defined interval. Set the timeout value between 100-6000 milliseconds. The default value is 1000 milliseconds.

5. Refer to the **Advanced RADIUS Configuration** field to set the WLAN's NAS configuration and RADIUS Dynamic Authorization.

NAS Identifier	Specify what should be included in the RADIUS NAS-Identifier field for authentication and accounting packets relating to this WLAN. Configuring a value here is optional, and defaults are used if this is not configured per WLAN.
NAS Port	The profile database on the RADIUS server consists of user profiles for each connected <i>network access server</i> (NAS) port. Each profile is matched to a username representing a physical port. When the wireless controller authorizes users, it queries the user profile database using a username representative of the physical NAS port making the connection. Set the numeric port value between 0-4,294,967,295.
RADIUS Dynamic Authorization	Select the check box to enable a mechanism that extends the RADIUS protocol to support unsolicited messages sent from the RADIUS server. These messages allow wireless network administrators to issue <i>change of authorization</i> (CoA) messages, which affect session authorization, or <i>Disconnect Message</i> (DM), which cause a session to be terminated immediately. This feature is disabled by default.

6. Refer to the **Radio Rates** field to define selected data rates for both the 2.4 and 5.0 GHz bands.

Rate Settings 2.4GHz-wlan

Radio Transmission Data Rates

b-only rates
 bg rates
 bgn rates
 Default
 g-only rates
 gn rates
 Custom Rates

802.11 b rates

	1Mbps	2Mbps	5.5Mbps	11Mbps
Basic:	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Supported:	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

802.11 g rates

	6Mbps	9Mbps	12Mbps	18Mbps	24Mbps	36Mbps	48Mbps	54Mbps
Basic:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Supported:	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

802.11 n rates

	MCS0-7	MCS8-15
Basic:	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Supported:	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

OK Reset Cancel

FIGURE 197 Advanced WLAN Rate Settings 2.4 GHz screen

FIGURE 198 Advanced WLAN Rate Settings 5 GHz screen

Define both minimum Basic and optimal Supported rates as required for the 802.11b rates, 802.11g rates and 802.11n rates supported by the 2.4 GHz band and 802.11a and 802.11n rates supported by the 5.0 GHz radio band. These are the rates wireless client traffic is supported within this WLAN.

If supporting 802.11n, select a Supported MCS index. Set a MCS (modulation and coding scheme) in respect to the radio's channel width and guard interval. A MCS defines (based on RF channel conditions) an optimal combination of 8 data rates, bonded channels, multiple spatial streams, different guard intervals and modulation types. Clients can associate as long as they support basic MCS (as well as non-11n basic rates).

The selected rates apply to associated client traffic within this WLAN only.

7. Select **OK** when completed to update this WLAN's advanced settings. Select **Reset** to revert the screen back to its last saved configuration.

Configuring WLAN QoS Policies

Wireless LAN Policy

QoS provides a data traffic prioritization scheme. QoS reduces congestion from excessive traffic. If there is enough bandwidth for all users and applications (unlikely because excessive bandwidth comes at a very high cost), then applying QoS has very little value. QoS provides policy enforcement for mission-critical applications and/or users that have critical bandwidth requirements when the wireless controller's bandwidth is shared by different users and applications.

QoS helps ensure each WLAN on the wireless controller receives a fair share of the overall bandwidth, either equally or as per the proportion configured. Packets directed towards clients are classified into categories such as Video, Voice and Data. Packets within each category are processed based on the weights defined for each WLAN.

The Quality of Service screen displays a list of QoS policies available to wireless controller managed WLANs. Each QoS policy has its own radio button that can be selected to edit its properties. If none of the existing QoS policies supports an ideal QoS configuration for the intended data traffic of this WLAN, select the Add button to create new policy. Select the radio button of an existing WLAN and select Ok to map the QoS policy to the WLAN displayed in the banner of the screen.

Use the WLAN *Quality of Service* (QoS) Policy screen to add a new QoS policy or edit the attributes of an existing policy.

NOTE

WLAN QoS configurations differ significantly from QoS policies configured for radios. WLAN QoS configurations are designed to support the data requirements of wireless clients, including the data types they support and their wireless controller managed network permissions. Radio QoS policies are specific to the transmit and receive characteristics of the connected radio's themselves, independent from the wireless clients these access point radios support.

1. Select **Configuration > Wireless > WLAN QoS Policy** to display existing QoS policies available to WLANs.

WLAN QoS Policy	Wireless Client Classification	SVP Prioritization	WMM Power Save	Multicast Mask Primary	Multicast Mask Secondary
default	VMM	✓	✓	11-22-AA-BB-AA-AA	22-11-AA-BB-AA-AA
WLANQoS1	VMM	✓	✓	11-22-AA-BB-AA-AA/FF-FF-F	22-11-AA-BB-AA-AA/FF-FF-F
WLANQoS2	VMM	✓	✓	11-22-AA-BB-AA-AA/AA-AB	22-11-AA-BB-AA-AA/AA-AB
WLANQoS3	VMM	✓	✓	11-22-AA-BB-AA-AA	22-11-AA-BB-AA-AA
WLANQoS4	VMM	✓	✓	11-22-AA-BB-AA-AA	22-11-AA-BB-AA-AA
WLANQoS5	VMM	✓	✓	11-22-AA-BB-AA-AA	22-11-AA-BB-AA-AA
WLANQoS6	VMM	✓	✓	11-22-AA-BB-AA-AA	22-11-AA-BB-AA-AA
WLANQoS7	VMM	✓	✓	11-22-AA-BB-AA-AA	22-11-AA-BB-AA-AA
WLANQoS8	VMM	✓	✓	11-22-AA-BB-AA-AA	22-11-AA-BB-AA-AA
WLANQoS9	VMM	✓	✓	11-22-AA-BB-AA-AA	22-11-AA-BB-AA-AA

Type to search in tables Row Count: 10

Add **Edit** **Delete**

FIGURE 199 WLAN Quality of Service (QoS) screen

2. Refer to the following read-only information on each listed QoS policy to determine whether an existing policy can be used as is, an existing policy requires edit or a new policy requires creation:

WLAN QoS Policy	Displays the name assigned to this WLAN QoS policy when it was initially created. The assigned policy name cannot be modified as part of the edit process.
Wireless Client Classification	<p>Lists each policy's Wireless Client Classification as defined for this WLAN's intended traffic. The Classification Categories are the different WLAN-WMM options available to a radio. Classification types include:</p> <p><i>WMM</i> – Implies WiFi Multimedia QoS extensions are enabled on this radio. This allows different traffic streams between the wireless client and the access point to be prioritized according to the type of traffic (voice, video etc). The WMM classification is required to support the high throughput data rates required of 802.11n device support.</p> <p><i>Voice</i> – Optimized for voice traffic. Implies all traffic on this WLAN is prioritized as voice traffic on the radio.</p> <p><i>Video</i> – Optimized for video traffic. Implies all traffic on this WLAN is prioritized as video traffic on the radio.</p> <p><i>Normal</i> – Optimized for best effort traffic. Implies all traffic on this WLAN is prioritized as best effort traffic on the radio.</p> <p><i>Low</i> – Optimized for background traffic. Implies all traffic on this WLAN is low priority on the radio.</p> <p><i>Non-Unicast</i> – Optimized for non-Unicast traffic. Implies all traffic on this WLAN is designed for broadcast or multicast.</p>
SVP Prioritization	A green checkmark defines the policy as having <i>Spectralink Voice Prioritization (SVP)</i> enabled to allow the wireless controller to identify and prioritize traffic from Spectralink/Polycomm phones using the SVP protocol. Phones using regular WMM and SIP are not impacted by SVP prioritization. A red "X" defines the QoS policy as not supporting SVP prioritization.
WMM Power Save	Enables support for the WMM based power-save mechanism, also known as <i>Unscheduled Automatic Power Save Delivery (U-APSD)</i> . This is primarily used by voice devices that are WMM capable. The default setting is enabled.
Multicast Mask Primary	Displays the primary multicast mask defined for each listed QoS policy. Normally all multicast and broadcast packets are buffered until the periodic DTIM interval (indicated in the 802.11 beacon frame), when clients in power save mode wake to check for frames. However, for certain applications and traffic types, the administrator may want the frames transmitted immediately, without waiting for the DTIM interval. By configuring a primary and secondary multicast mask, an administrator can indicate which frames are transmitted immediately. Setting masks is optional and only needed if there are traffic types requiring special handling.
Multicast Mask Secondary	Displays the secondary multicast mask defined for each listed QoS policy.

NOTE

When using a wireless client classification other than WMM, only legacy rates are supported on that WLAN.

3. Either select the **Add** button to define a new WLAN QoS policy, or select an existing WLAN QoS policy and select **Edit** to modify its existing configuration. Existing QoS policies can be selected and deleted as needed.

A *Quality of Service (QoS)* policy screen displays for the new or selected WLAN. The screen displays the WMM tab by default, but additional tabs also display for WLAN and wireless client rate limit configurations. For more information, refer to the following:

- [Configuring a WLAN's QoS WMM Settings](#)
- [Configuring Rate Limit Settings](#)
- [Configuring Multimedia Optimizations](#)

Configuring a WLAN's QoS WMM Settings

Using WMM, end-user satisfaction is maintained in a wider variety of environments and traffic conditions. WMM makes it possible for both home networks and Enterprises to decide which data streams are most important and assign them a higher traffic priority.

WMM's prioritization capabilities are based on the four access categories. The higher the access category, the higher the probability to transmit this kind of traffic over the wireless controller managed WLAN. ACs were designed to correspond to 802.1d priorities to facilitate interoperability with QoS policy management mechanisms. WMM enabled wireless controllers/ access points coexist with legacy devices (not WMM-enabled).

Packets not assigned to a specific access category are categorized by default as having best effort priority. Applications assign each data packet to a given access category packets are then added to one of four independent transmit queues (one per access category - voice, video, best effort or background) in the client. The client has an internal collision resolution mechanism to address collision among different queues, which selects the frames with the highest priority to transmit.

The same mechanism deals with external collision, to determine which client should be granted the *opportunity to transmit* (TXOP). The collision resolution algorithm responsible for traffic prioritization is probabilistic and depends on two timing parameters that vary for each access category.

- The minimum interframe space, or *Arbitrary Inter-Frame Space Number* (AIFSN)
- The contention window, sometimes referred to as the random backoff wait

Both values are smaller for high-priority traffic. The value of the contention window varies through time. Initially the contention window is set to a value that depends on the AC. As frames with the highest AC tend to have the lowest backoff values, they are more likely to get a TXOP.

After each collision the contention window is doubled until a maximum value (also dependent on the AC) is reached. After successful transmission, the contention window is reset to its initial, AC dependant value. The AC with the lowest backoff value gets the TXOP.

To configure a WMM configuration for a wireless controller managed WLAN:

1. Select **Configuration > Wireless > WLAN QoS Policy** to display existing QoS Policies available to the wireless controller managed network.
2. Select the **Add** button to create a new QoS policy or **Edit** to modify the properties of an existing WLAN QoS policy.

The WMM tab displays by default.

WLAN QoS Policy WLANQoS2

WMM Rate Limit Multimedia Optimizations

Settings

- Wireless Client Classification: vWMM
- Non-Unicast Classification: Normal
- Enable Voice Prioritization:
- Enable SVP Prioritization:
- Enable WMM Power Save:
- Enable GBSS Load IE:
- Configure Non WMM Client Traffic: Normal

Voice Access

- Transmit Ops: 47 (0 to 65,535)
- AIFS N: 2 (2 to 15)
- ECW Min: 2 (0 to 15)
- ECW Max: 3 (0 to 15)

Normal (Best Effort) Access

- Transmit Ops: 25 (0 to 65,535)
- AIFS N: 3 (2 to 15)
- ECW Min: 4 (0 to 15)
- ECW Max: 10 (0 to 15)

Video Access

- Transmit Ops: 94 (0 to 65,535)
- AIFS N: 2 (2 to 15)
- ECW Min: 3 (0 to 15)
- ECW Max: 4 (0 to 15)

Low (Background) Access

- Transmit Ops: 25 (0 to 65,535)
- AIFS N: 7 (2 to 15)
- ECW Min: 4 (0 to 15)
- ECW Max: 10 (0 to 15)

Other Settings

- Trust IP DSCP:
- Trust 802.11 WMM QoS:

OK Reset Exit

FIGURE 200 WLAN QoS Policy WMM screen

3. Configure the following in respect to the WLAN's intended WMM radio traffic and user requirements:

Wireless Client Classification

Use the drop-down menu to select the Wireless Client Classification for this WLAN's intended traffic. The Classification Categories are the different WLAN-WMM options available to the radio. The Wireless Client Classification types are:

WMM – Implies WiFi Multimedia QoS extensions are enabled on this radio. This allows different traffic streams between the wireless client and the access point to be prioritized according to the type of traffic (voice, video etc). The WMM classification is required to support the high throughput data rates required of 802.11n device support.

Voice – Optimized for voice traffic. Implies all traffic on this WLAN is prioritized as voice traffic on the radio.

Video – Optimized for video traffic. Implies all traffic on this WLAN is prioritized as video traffic on the radio.

Normal – Optimized for best effort traffic. Implies all traffic on this WLAN is prioritized as best effort traffic on the radio.

Low – Optimized for background traffic. Implies all traffic on this WLAN is low priority on the radio.

Non-Unicast – Optimized for non-Unicast traffic. Implies all traffic on this WLAN is designed for broadcast or multiple destinations.

Non-Unicast Classification

Use the drop-down menu to select the Non-Unicast Classification for this WLAN's intended traffic. The Non-Unicast Classification types are:

Voice – Optimized for voice traffic. Implies all traffic on this WLAN is prioritized as voice traffic on the radio.

Video – Optimized for video traffic. Implies all traffic on this WLAN is prioritized as video traffic on the radio.

Normal – Optimized for best effort traffic. Implies all traffic on this WLAN is prioritized as best effort traffic on the radio.

Low – Optimized for background traffic. Implies all traffic on this WLAN is low priority on the radio.

Enable Voice Prioritization

Select this option if Voice traffic is prioritized on the WLAN. This gives priority to voice and voice management packets and is supported only on certain legacy Brocade Mobility VOIP phones. This feature is disabled by default.

Enable SVP Prioritization

Enabling *Spectralink Voice Prioritization (SVP)* allows the wireless controller to identify and prioritize traffic from Spectralink/Polycomm phones. This gives priority to voice, with voice management packets supported only on certain legacy Brocade Mobility VOIP phones. If the Wireless Client Classification is WMM, non WMM devices recognized as voice devices have all their traffic transmitted at voice priority. Devices are classified as voice, when they emit SIP, SCCP, or H323 traffic. Thus, selecting this option has no effect on devices supporting WMM. This feature is disabled by default.

Enable WMM Power Save	Enables support for the WMM based power-save mechanism, also known as <i>Unscheduled Automatic Power Save Delivery (U-APSD)</i> . This is primarily used by voice devices that are WMM capable. The default setting is enabled.
Enable QBSS Load IE	Check this box to enable QoS Basis Service Set (QBSS) information element (IE) in beacons and probe response packets advertised by access points. The default value is enabled.
Configure Non WMM Client Traffic	Use the drop-down menu to select the Non-WMM client traffic Classification. The Non-WMM Classification types are: <i>Voice</i> – Optimized for voice traffic. Implies all traffic on this WLAN is prioritized as voice traffic on the radio. <i>Video</i> – Optimized for video traffic. Implies all traffic on this WLAN is prioritized as video traffic on the radio. <i>Normal</i> – Optimized for best effort traffic. Implies all traffic on this WLAN is prioritized as best effort traffic on the radio. <i>Low</i> – Optimized for background traffic. Implies all traffic on this WLAN is low priority on the radio.

4. Set the following **Voice Access** settings for the WLAN's QoS policy:

Transmit Ops	Use the slider to set the maximum duration a device can transmit after obtaining a transmit opportunity. The default value is 47.
AIFSN	Set the current <i>Arbitrary Inter-frame Space Number (AIFSN)</i> between 2-15. Higher-priority traffic voice categories should have lower AIFSNs than lower-priority traffic categories. This will cause lower-priority traffic to wait longer before attempting access. The default value is 2.
ECW Min	The ECW Min is combined with the ECW Max to create the contention value in the form of a numerical range. From this range, a random number is selected for the back off mechanism. Lower values are used for higher priority traffic. The available range to the wireless controller is from 0-15. The default value is 2.
ECW Max	The ECW Max is combined with the ECW Min to create the contention value in the form of a numerical range. From this range, a random number is selected for the back off mechanism. Lower values are used for higher priority traffic. The available range to the wireless controller is from 0-15. The default value is 3.

5. Set the following **Normal (Best Effort) Access** settings for the WLAN's QoS policy:

Transmit Ops	Use the slider to set the maximum duration a device can transmit after obtaining a transmit opportunity. The default value is 0.
AIFSN	Set the current AIFSN between 2-15. The default value is 3.
ECW Min	The ECW Min is combined with the ECW Max to create the contention value in the form of a numerical range. From this range, a random number is selected for the back off mechanism. Higher values are used for lower priority traffic (like Normal). The available range to the wireless controller is from 0-15. The default value is 4.
ECW Max	The ECW Max is combined with the ECW Min to create the contention value in the form of a numerical range. From this range, a random number is selected for the back off mechanism. Higher values are used for lower priority traffic (like Normal). The available range to the wireless controller is from 0-15. The default value is 10.

6. Set the following **Video Access** settings for the WLAN's QoS policy:

Transmit Ops	Use the slider to set the maximum duration a device can transmit after obtaining a transmit opportunity. The default values is 94.
AIFS	Set the current <i>Arbitrary Inter-frame Space Number</i> (AIFS) between 2-15. Higher-priority traffic video categories should have lower AIFSNs than lower-priority traffic categories. This will cause lower-priority traffic to wait longer before attempting access. The default value is 2.
ECW Min	The ECW Min is combined with the ECW Max to create the contention value in the form of a numerical range. From this range, a random number is selected for the back off mechanism. Lower values are used for higher priority traffic (like video). The available range to the wireless controller is from 0-15. The default value is 3.
ECW Max	The ECW Max is combined with the ECW Min to create the contention value in the form of a numerical range. From this range, a random number is selected for the back off mechanism. Lower values are used for higher priority traffic (like video). The available range to the wireless controller is from 0-15. The default value is 4.

7. Set the following **Low (Background) Access** settings for the WLAN's QoS policy:

Transmit Ops	Use the slider to set the maximum duration a device can transmit after obtaining a transmit opportunity. For higher-priority traffic categories, this value should be set to a low number. The default value is 0.
AIFS	Set the current AIFS between 2-15. The default value is 7.
ECW Min	The ECW Min is combined with the ECW Max to create the contention value in the form of a numerical range. From this range, a random number is selected for the back off mechanism. Higher values are used for lower priority traffic (like Low). The available range to the wireless controller is from 0-15. The default value is 4.
ECW Max	The ECW Max is combined with the ECW Min to create the contention value in the form of a numerical range. From this range, a random number is selected for the back off mechanism. Higher values are used for lower priority traffic (like Low). The available range to the wireless controller is from 0-15. The default value is 10.

8. Set the following **Other Settings** for the WLAN's QoS policy:

Trust IP DSCP	Select this option to trust IP DSCP values for WLANs. The default value is enabled.
Trust 802.11 WMM QoS	Select this option to trust 802.11 WMM QoS values for WLANs. The default value enabled.

9. Select **OK** when completed to update this WLAN's QoS settings. Select **Reset** to revert the screen back to its last saved configuration.

Configuring Rate Limit Settings

Excessive traffic can cause performance issues or bring down the network entirely. Excessive traffic can be caused by numerous sources including network loops, faulty devices or malicious software such as a worm or virus that has infected on one or more devices at the branch. Rate limiting limits the maximum rate sent to or received from the wireless network (and WLAN) per wireless client. It prevents any single user from overwhelming the wireless network. It can also provide differential service for service providers. The uplink and downlink rate limits are usually configured on a RADIUS server using Brocade Mobility vendor specific attributes. The controller extracts the rate limits from RADIUS server's response. When such attributes are not present, the

settings defined on the controller are applied. An administrator can set separate QoS rate limit configurations for data transmitted from the managed network (upstream) and data transmitted from a WLAN's wireless clients back to their associated access point radios and controller (downstream).

Before defining rate limit thresholds for WLAN upstream and downstream traffic, Brocade Mobility recommends you define the normal number of ARP, broadcast, multicast and unknown unicast packets that typically transmit and receive from each supported WMM access category. If thresholds are defined too low, normal network traffic (required by end-user devices) will be dropped by the controller resulting in intermittent outages and performance problems.

A controller's connected wireless clients can also have QoS rate limit settings defined in both the upstream and downstream direction.

To configure a QoS rate limit configuration for a managed WLAN:

1. Select **Configuration > Wireless > WLAN QoS Policy** to display existing QoS policies available to WLANs.
2. Either select the **Add** button to define a new WLAN QoS policy, or select an existing WLAN QoS policy and select **Edit** to modify its existing configuration.
3. Select the **Rate Limit** tab.

The screenshot displays the 'Rate Limit' configuration screen for a WLAN QoS policy. The interface is divided into two main sections, each with a 'Downstream Rate Limit' and a 'Downstream Random Early Detection Threshold'.

Top Section (Disabled):

- Downstream Rate Limit:**
 - Enable:
 - Rate: 5000 (50 to 1,000,000 kbps)
 - Maximum Burst Size: 320 (2 to 1,024 kbytes)
- Downstream Random Early Detection Threshold:**
 - Background Traffic: 50 (0 to 100 %)
 - Best Effort Traffic: 50 (0 to 100 %)
 - Video Traffic: 25 (0 to 100 %)
 - Voice Traffic: 0 (0 to 100 %)

Bottom Section (Enabled):

- Downstream Rate Limit:**
 - Enable:
 - Rate: 1000 (50 to 1,000,000 kbps)
 - Maximum Burst Size: 64 (2 to 1,024 kbytes)
- Downstream Random Early Detection Threshold:**
 - Background Traffic: 50 (0 to 100 %)
 - Best Effort Traffic: 50 (0 to 100 %)
 - Video Traffic: 25 (0 to 100 %)
 - Voice Traffic: 0 (0 to 100 %)

At the bottom right, there are buttons for 'OK', 'Reset', and 'Exit'.

FIGURE 201 QoS Policy WLAN Rate Limit screen

4. Configure the following parameters in respect to the intended **WLAN Upstream Rate Limit**, or traffic from the controller to associated access point radios and their associated wireless clients:

Enable	Select the Enable check box to enable rate limiting for data transmitted from the controller to its connected access point radios and associated wireless clients. Enabling this option does not invoke rate limiting for data traffic in the downstream direction. This feature is disabled by default.
Rate	Define an upstream rate limit between 50 - 1,000,000 kbps. This limit constitutes a threshold for the maximum the number of packets transmitted or received over the WLAN (from all access categories). Traffic that exceeds the defined rate is dropped by the controller and a log message is generated. The default setting is 5000 kbps.
Maximum Burst Size	Set a maximum burst size between 2 - 1024 kbytes. The smaller the burst, the less likely the upstream packet transmission will result in congestion for the WLAN's client destinations. By trending the typical number of ARP, broadcast, multicast and unknown unicast packets over a period of time, the average rate for each access category can be obtained. Once a baseline is obtained, administrators should then add a 10% margin (minimally) to allow for traffic bursts at the site. The default burst size is 64 kbytes.

5. Set the following **WLAN Upstream Random Early Detection Threshold** settings for each access category. An early random drop is done when a traffic stream falls below the set threshold.

Background Traffic	Set a percentage value for background traffic in the upstream direction. This is a percentage of the maximum burst size for low priority traffic. Background traffic exceeding the defined threshold is dropped by the controller and a log message is generated. Background traffic consumes the least bandwidth of any access category, so this value can be set to a lower value once a general upstream rate is known by the network administrator (using a time trend analysis). The default threshold is 50%.
Best Effort Traffic	Set a percentage value for best effort traffic in the upstream direction. This is a percentage of the maximum burst size for normal priority traffic. Best effort traffic exceeding the defined threshold is dropped by the controller and a log message is generated. Best effort traffic consumes little bandwidth, so this value can be set to a lower value once a general upstream rate is known by the network administrator (using a time trend analysis). The default threshold is 50%.
Video Traffic	Set a percentage value for video traffic in the upstream direction. This is a percentage of the maximum burst size for video traffic. Video traffic exceeding the defined threshold is dropped by the controller and a log message is generated. Video traffic consumes significant bandwidth, so this value can be set to a higher value once a general upstream rate is known by the network administrator (using a time trend analysis). The default threshold is 25%.
Voice Traffic	Set a percentage value for voice traffic in the upstream direction. This is a percentage of the maximum burst size for voice traffic. Voice traffic exceeding the defined threshold is dropped by the controller and a log message is generated. Voice applications consume significant bandwidth, so this value can be set to a higher value once a general upstream rate is known by the network administrator (using a time trend analysis). The default threshold is 0%.

6. Configure the following parameters in respect to the intended **WLAN Downstream Rate Limit**, or traffic from wireless clients to associated Access Point radios and the controller:

Enable	Select the Enable radio button to enable rate limiting for data transmitted from the controller to its connected access point radios and associated wireless clients. Enabling this option does not invoke rate limiting for data traffic in the upstream direction. This feature is disabled by default.
Rate	Define an upstream rate limit between 50 - 1,000,000 kbps. This limit constitutes a threshold for the maximum the number of packets transmitted or received over the WLAN (from all access categories). Traffic that exceeds the defined rate is dropped by the controller and a log message is generated. The default setting is 5000 kbps.
Maximum Burst Size	Set a maximum burst size between 2 - 1024 kbytes. The smaller the burst, the less likely the downstream packet transmission will result in congestion for the WLANs wireless client destinations. By trending the typical number of ARP, broadcast, multicast and unknown unicast packets over a period of time, the average rate for each access category can be obtained. Once a baseline is obtained, administrators should then add a minimum of a 10% margin to allow for traffic bursts at the site. The default burst size is 64 kbytes.

7. Set the following **WLAN Downstream Random Early Detection Threshold** settings for each access category. An early random drop is done when the amount of tokens for a traffic stream falls below the set threshold.

Background Traffic	Set a percentage value for background traffic in the downstream direction. This is a percentage of the maximum burst size for low priority traffic. Background traffic exceeding the defined threshold is dropped by the controller and a log message is generated. Background traffic consumes the least bandwidth of any access category, so this value can be set to a lower value once a general downstream rate is known by the network administrator (using a time trend analysis). The default threshold is 50%.
Best Effort Traffic	Set a percentage value for best effort traffic in the downstream direction. This is a percentage of the maximum burst size for normal traffic. Best effort traffic exceeding the defined threshold is dropped by the controller and a log message is generated. Best effort traffic consumes little bandwidth, so this value can be set to a lower value once a general downstream rate is known by the network administrator (using a time trend analysis). The default threshold is 50%.
Video Traffic	Set a percentage value for video traffic in the downstream direction. This is a percentage of the maximum burst size for video traffic. Video traffic exceeding the defined threshold is dropped by the controller and a log message is generated. Video traffic consumes significant bandwidth, so this value can be set to a higher value once a general downstream rate is known by the network administrator (using a time trend analysis). The default threshold is 25%.
Voice Traffic	Set a percentage value for voice traffic in the downstream direction. This is a percentage of the maximum burst size for voice traffic. Voice traffic exceeding the defined threshold is dropped by the controller and a log message is generated. Voice applications consume significant bandwidth, so this value can be set to a higher value once a general upstream rate is known by the network administrator (using a time trend analysis). The default threshold is 0%. 0% means no early random drops will occur.

8. Configure the following parameters in respect to the intended **Wireless Client Upstream Rate Limit**:

Enable	Select the Enable radio button to enable rate limiting for data transmitted from the client to its associated access point radio and connected wireless controller. Enabling this option does not invoke client rate limiting for data traffic in the downstream direction. This feature is disabled by default.
Rate	Define an upstream rate limit between 50 - 1,000,000 kbps. This limit constitutes a threshold for the maximum the number of packets transmitted or received (from all access categories). Traffic that exceeds the defined rate is dropped by the client and a log message is generated. The default rate is 1,000 kbps.
Maximum Burst Size	Set a maximum burst size between 2 - 1024 kbytes. The smaller the burst, the less likely the upstream packet transmission will result in congestion for the wireless client. The default burst size is 64 kbytes.

9. Set the following **Wireless Client Upstream Random Early Detection Threshold** settings for each access category:

Background Traffic	Set a percentage value for background traffic in the upstream direction. This is a percentage of the maximum burst size for low priority traffic. Background traffic exceeding the defined threshold is dropped by the client and a log message is generated. The default threshold is 50%.
Best Effort Traffic	Set a percentage value for best effort traffic in the upstream direction. This is a percentage of the maximum burst size for normal traffic. Best effort traffic exceeding the defined threshold is dropped by the client and a log message is generated. The default threshold is 50%.
Video Traffic	Set a percentage value for video traffic in the upstream direction. This is a percentage of the maximum burst size for video traffic. Video traffic exceeding the defined threshold is dropped by the client and a log message is generated. The default threshold is 25%.
Voice Traffic	Set a percentage value for voice traffic in the downstream direction. This is a percentage of the maximum burst size for voice traffic. Voice traffic exceeding the defined threshold is dropped by the client and a log message is generated. The default threshold is 0%. 0% implies no early random drops will occur.

10. Configure the following parameters in respect to the intended **Wireless Client Downstream Rate Limit**, or traffic from a controller to associated access point radios and the wireless client:

Enable	Select the Enable radio button to enable rate limiting for data transmitted from connected wireless clients to the controller. Enabling this option does not invoke rate limiting for data traffic in the upstream direction. This feature is disabled by default.
Rate	Define a downstream rate limit between 50 - 1,000,000 kbps. This limit constitutes a threshold for the maximum the number of packets transmitted or received by the client. Traffic that exceeds the defined rate is dropped and a log message is generated. The default rate is 1,000 kbytes.
Maximum Burst Size	Set a maximum burst size between 2 - 64 kbytes. The smaller the burst, the less likely the downstream packet transmission will result in congestion for the wireless client. The default burst size is 6 kbytes.

11. Set the following **Wireless Client sDownstream Random Early Detection Threshold** settings for each access category:

Background Traffic	Set a percentage value for background traffic in the downstream direction. This is a percentage of the maximum burst size for low priority traffic. Background traffic exceeding the defined threshold is dropped by the client and a log message is generated. The default is 50%.
Best Effort Traffic	Set a percentage value for best effort traffic in the downstream direction. This is a percentage of the maximum burst size for normal traffic. Best effort traffic exceeding the defined threshold is dropped by the client and a log message is generated. The default is 50%.
Video Traffic	Set a percentage value for video traffic in the downstream direction. This is a percentage of the maximum burst size for video traffic. Video traffic exceeding the defined threshold is dropped by the client and a log message is generated. The default is 25%.
Voice Traffic	Set a percentage value for voice traffic in the downstream direction. This is a percentage of the maximum burst size for voice traffic. Voice traffic exceeding the defined threshold is dropped by the client and a log message is generated. The default threshold is 0%.0% means no early random drops will occur.

12. Select **OK** when completed to update this WLAN's QoS rate limit settings. Select **Reset** to revert the screen back to its last saved configuration.

WLAN QoS Deployment Considerations

Before defining a QoS configuration on a wireless controller WLAN, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- WLAN QoS configurations differ significantly from QoS policies configured for wireless controller associated access point radios. WLAN QoS configurations are designed to support the data requirements of wireless clients, including the data types they support and their wireless controller managed network permissions. Radio QoS policies are specific to the transmit and receive characteristics of the connected radio's themselves, independent from the wireless clients these access point radios support.
- Enabling WMM support on a wireless controller's WLAN only advertises WMM capability to wireless clients. The wireless clients must be also able to support WMM and use the parameters correctly while accessing the wireless network to truly benefit.
- Rate limiting is disabled by default on all controller managed WLANs. To enable rate limiting, a threshold must be defined for WLAN.
- Before enabling rate limiting on a controller managed WLAN, a baseline for each traffic type should be performed. Once a baseline has been determined, a minimum 10% margin should be added to allow for traffic bursts.
- The bandwidth required for real-time applications such as voice and video are very fairly easy to calculate as the bandwidth requirements are consistent and can be realistically trended over time. Applications such as Web, database and email are harder to estimate, since bandwidth usage varies depending on how the applications are utilized.

Configuring Multimedia Optimizations

To configure multimedia optimizations for a controller managed WLAN:

1. Select **Configuration > Wireless > WLAN QoS Policy** to display existing QoS policies available to WLANs.
2. Either select the **Add** button to define a new WLAN QoS policy, or select an existing WLAN QoS policy and select **Edit** to modify its existing configuration.
3. Select the **Multimedia Optimizations** tab.

FIGURE 202 QoS Policy Multimedia Optimizations screen

4. Configure the following parameters in respect to the intended **Multicast Mask**:

Multicast Mask Primary

Configure the primary multicast mask defined for each listed QoS policy. Normally all multicast and broadcast packets are buffered until the periodic DTIM interval (indicated in the 802.11 beacon frame), when clients in power save mode wake to check for frames. However, for certain applications and traffic types, an administrator may want the frames transmitted immediately, without waiting for the DTIM interval. By configuring a primary and secondary multicast mask, an administrator can indicate which frames are transmitted immediately. Setting masks is optional and only needed if there are traffic types requiring special handling.

Multicast Mask Secondary

Set a secondary multicast mask for the WLAN QoS policy. Normally all multicast and broadcast packets are buffered until the periodic DTIM interval (indicated in the 802.11 beacon frame), when clients in power save mode wake to check for frames. However, for certain applications and traffic types, an administrator may want the frames transmitted immediately, without waiting for the DTIM interval. By configuring a primary and secondary multicast mask, an administrator can indicate which frames are transmitted immediately. Setting masks is optional and only needed if there are traffic types requiring special handling.

5. Set the following **Accelerated Multicast** settings:

Disable Multicast Streaming	Select this option to disable all Multicast Streaming on the WLAN.
Automatically Detect Multicast Streams	Select this option to allow the administrator to have multicast packets that are being bridged converted to unicast to provide better overall airtime utilization and performance. The administrator can either have the system automatically detect multicast streams and convert all detected multicast streams to unicast, or specify which multicast streams are to be converted to unicast. When the stream is converted and being queued up for transmission, there are a number of classification mechanisms that can be applied to the stream and the administrator can select what type of classification they would want.
Manually Configure Multicast Addresses	Select this option and specify a list of multicast addresses and classifications. Packets are accelerated when the destination addresses matches.

6. Select **OK** when completed to update this WLAN's Multimedia Optimizations settings. Select **Reset** to revert the screen back to its last saved configuration.

WLAN QoS Deployment Considerations

Before defining a QoS configuration on a wireless controller WLAN, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- WLAN QoS configurations differ significantly from QoS policies configured for wireless controller associated access point radios. WLAN QoS configurations are designed to support the data requirements of wireless clients, including the data types they support and their wireless controller managed network permissions. Radio QoS policies are specific to the transmit and receive characteristics of the connected radio's themselves, independent from the wireless clients these access point radios support.
- Enabling WMM support on a wireless controller's WLAN only advertises WMM capability to wireless clients. The wireless clients must be also able to support WMM and use the parameters correctly while accessing the wireless network to truly benefit.
- Rate limiting is disabled by default on all controller managed WLANs. To enable rate limiting, a threshold must be defined for WLAN.
- Before enabling rate limiting on a controller managed WLAN, a baseline for each traffic type should be performed. Once a baseline has been determined, a minimum 10% margin should be added to allow for traffic bursts.
- The bandwidth required for real-time applications such as voice and video are very fairly easy to calculate as the bandwidth requirements are consistent and can be realistically trended over time. Applications such as Web, database and email are harder to estimate, since bandwidth usage varies depending on how the applications are utilized.

Radio QoS Policy

Without a dedicated QoS policy, a controller managed network operates on a best-effort delivery basis, meaning all traffic has equal priority and an equal chance of being delivered in a timely manner. When congestion occurs, all traffic has an equal chance of being dropped!

When configuring a QoS policy for a controller radio, select specific network traffic, prioritize it, and use congestion-management and congestion-avoidance techniques to provide deployment customizations best suited to each QoS policy's intended wireless client base.

Brocade Mobility controllers and their associated access point radios and wireless clients support several *Quality of Service* (QoS) techniques enabling real-time applications (such as voice and video) to co-exist simultaneously with lower priority background applications (such as Web, Email and file transfers). A well designed QoS policy should:

- Classify and mark data traffic to accurately prioritize and segregate it (by access category) throughout the controller managed network.
- Minimize the network delay and jitter for latency sensitive traffic.
- Ensure higher priority traffic has a better likelihood of delivery in the event of network congestion.
- Prevent the ineffective utilization of access points degrading session quality by configuring admission control mechanisms within each radio QoS policy

Within a Brocade Mobility controller managed network, wireless clients supporting low and high priority traffic contend with one another for controller access and data resources. The IEEE 802.11e amendment has defined *Enhanced Distributed Channel Access* (EDCA) mechanisms stating high priority traffic can access the controller managed network sooner than lower priority traffic. The EDCA defines four traffic classes (or access categories); voice (highest), video (next highest), best effort and background (lowest). The EDCA has defined a time interval for each traffic class, known as the *Transmit Opportunity* (TXOP). The TXOP prevents traffic of a higher priority from completely dominating the wireless medium, thus ensuring lower priority traffic is still supported by the controller and its connected radios.

IEEE 802.11e includes an advanced power saving technique called *Unscheduled Automatic Power Save Delivery* (U-APSD) that provides a mechanism for wireless clients to retrieve packets buffered by an access point. U-APSD reduces the amount of signaling frames sent from a client to retrieve buffered data from an access point. U-APSD also allows access points to deliver buffered data frames as *bursts*, without backing-off between data frames. These improvements are useful for voice clients, as they provide improved battery life and call quality.

The Wi-Fi alliance has created *Wireless Multimedia* (WMM) and *WMM Power Save* (WMM-PS) certification programs to ensure interoperability between 802.11e WLAN infrastructure implementations and wireless clients. A Brocade Mobility controller managed wireless network supports both WMM and WMM-Power Save techniques. WMM and WMM-PS (U-APSD) are enabled by default in each controller WLAN profile.

Enabling WMM support on a controller managed WLAN just advertises the WLAN's WMM capability and radio configuration to wireless clients. The wireless clients must be also able to support WMM and use the values correctly while accessing the controller's WLAN to benefit.

WMM includes advanced parameters (CWMin, CWMax, AIFSN and TXOP) specifying back-off duration and inter-frame spacing when accessing the controller managed network. These parameters are relevant to both connected access point radios and their wireless clients. Parameters impacting access point transmissions to their clients are controlled using per radio WMM settings, while parameters used by wireless clients are controlled by a WLAN's WMM settings.

Brocade Mobility RFS4000, RFS6000 and RFS7000 model controllers include a *Session Initiation Protocol (SIP)*, *Skinny Call Control Protocol (SCCP)* and *Application Layer Gateway (ALGs)* enabling devices to identify voice streams and dynamically set voice call bandwidth. Controllers use the data to provide prioritization and admission control to these devices without requiring TSPEC or WMM client support.

Brocade Mobility controllers support static QoS mechanisms per WLAN to provide prioritization of WLAN traffic when legacy (non WMM) clients are deployed. A controller and AP-650 access point allow flexible WLAN mapping with a static WMM access control value. When enabled on a WLAN, traffic forwarded from a controller to a client is prioritized and forwarded based on the WLAN's WMM access control setting.

NOTE

Statically setting a WLAN WMM access category value only prioritizes traffic from the controller to the client.

Wireless network administrators can also assign weights to each WLAN in relation to user priority levels. The lower the weight, the lower the priority. Use a weighted round robin technique to achieve different QoS levels across the WLANs supported by the controller.

The controller provides the means to rate-limit bandwidth for WLAN sessions. This form of per-user rate limiting enables administrators to define uplink and downlink bandwidth limits for users and clients. This sets the level of traffic a user or client can forward and receive over the WLAN. If the user or client exceeds the limit, the controller drops the excessive traffic.

Rate limits can be applied to WLANs using groups defined on the controller or externally from a RADIUS server using Brocade Mobility *Vendor Specific Attributes (VSAs)*. Rate limits can be applied to users authenticating to the controller using 802.1X, hotspot authentication and devices using MAC authentication.

Configuring Radio QoS Policies

[Radio QoS Policy](#)

To configure a radio's QoS policy:

1. Select **Configuration > Wireless > Radio QoS Policy** to display existing Radio QoS policies.

Best Effort

A green checkmark indicates that Best Effort QoS is enabled on the radio. A red X indicates that Best Effort QoS is disabled on the radio.

Video

A green checkmark indicates that Video prioritization QoS is enabled on the radio. A red X indicates that Video prioritization QoS is disabled on the radio.

Background

A green checkmark indicates that Background prioritization QoS is enabled on the radio. A red X indicates that Background prioritization QoS is disabled on the radio.

3. Either select **Add** to create a new radio QoS policy, or select one of the existing policies listed and select the **Edit** button to modify its configuration.

The screenshot displays the 'WMM' tab of the Radio QoS Policy configuration screen. It is divided into four sections, each with a 'Transmit Ops' field and three 'ECW' (AIFS, Min, Max) fields. The 'Normal (Best Effort) Access' section is currently selected.

Access Category	Transmit Ops	ECW Min	ECW Max
Voice Access	47	1	3
Video Access	94	1	4
Normal (Best Effort) Access	31	3	6
Low (Background) Access	0	7	10

At the bottom right of the screen are buttons for 'OK', 'Reset', and 'Exit'.

FIGURE 204 Radio QoS Policy WMM screen

The Radio QoS Policy screen displays the **WMM** tab by default. Use the WMM tab to define the access category configuration (*CWMin*, *CWMax*, *AIFS*N and *TXOP* values) in respect to the type of wireless data planned for this new or updated WLAN radio QoS policy.

4. Set the following **Voice Access** settings for the Radio QoS policy:

Transmit Ops	Use the slider to set the maximum duration a device can transmit after obtaining a transmit opportunity. When controller resources are shared between a <i>Voice over IP</i> (VoIP) call and a low priority file transfer, bandwidth is normally exploited by the file transfer, thus reducing call quality or even causing the call to disconnect. With voice QoS, a VoIP call (a real-time session), receives priority, maintaining a high level of voice quality. For higher-priority traffic categories (like voice), the Transmit Ops value should be set to a low number. The default value is 47.
AIFSN	Set the current AIFSN between 1-15. Higher-priority traffic voice categories should have lower AIFSNs than lower-priority traffic categories. This will cause lower-priority traffic to wait longer before attempting access. The default value is 1.
ECW Min	The ECW Min is combined with the ECW Max to create a contention value in the form of a numerical range. From this range, a random number is selected for the back off mechanism. Lower values are used for higher priority traffic. The available range to the controller is from 0-15. The default value is 2.
Power Save	The ECW Max is combined with the ECW Min to create a contention value in the form of a numerical range. From this range, a random number is selected for the back off mechanism. Lower values are used for higher priority traffic. The available range to the controller is from 0-15. The default value is 3.

5. Set the following **Normal (Best Effort) Access** settings for the radio QoS policy:

Transmit Ops	Use the slider to set the maximum duration a device can transmit after obtaining a transmit opportunity. For higher-priority traffic categories, this value should be set to a low number. The default value is 0.
AIFSN	Set the current AIFSN between 1-15. Higher-priority traffic voice categories should have lower AIFSNs than lower-priority traffic categories. This will cause lower-priority traffic to wait longer before attempting access. The default value is 7.
ECW Min	The ECW Min is combined with the ECW Max to create a contention value in the form of a numerical range. From this range, a random number is selected for the back off mechanism. Higher values are used for lower priority traffic (like Normal). The available range to the controller is from 0-15. The default value is 4.
Power Save	The ECW Max is combined with the ECW Min to create a contention value in the form of a numerical range. From this range, a random number is selected for the back off mechanism. Higher values are used for lower priority traffic (like Normal). The available range to the controller is from 0-15. The default value is 10.

6. Set the following **Video Access** settings for the Radio QoS policy:

Transmit Ops	Use the spinner control to set the maximum duration a radio can transmit after obtaining a transmit opportunity. For higher-priority traffic categories (like video), this value should be set to a low number. The default value is 94.
---------------------	--

AIFSN	Set the current AIFSN between 1-15. Higher-priority traffic video categories should have lower AIFSNs than lower-priority traffic categories. This will cause lower-priority traffic to wait longer before attempting access. The default value is 1.
ECW Min	The ECW Min is combined with the ECW Max to create a contention value in the form of a numerical range. From this range, a random number is selected for the back off mechanism. Lower values are used for higher priority traffic (like video). The available range to the controller is from 0-15. The default value is 3.
ECW Max	The ECW Max is combined with the ECW Min to create a contention value in the form of a numerical range. From this range, a random number is selected for the back off mechanism. Lower values are used for higher priority traffic (like video). The available range to the controller is from 0-15. The default value is 4.

7. Set the following **Low (Best Effort) Access** settings for the radio QoS policy:

Transmit Ops	Use the slider to set the maximum duration a device can transmit after obtaining a transmit opportunity. For higher-priority traffic categories, this value should be set to a low number. The default value is 0.
AIFSN	Set the current AIFSN between 1-15. Higher-priority traffic voice categories should have lower AIFSNs than lower-priority traffic categories. This will cause lower-priority traffic to wait longer before attempting access. The default value is 3.
ECW Min	The ECW Min is combined with the ECW Max to create a contention value in the form of a numerical range. From this range, a random number is selected for the back off mechanism. Higher values are used for lower priority traffic (like Low). The available range to the controller is from 0-15. The default value is 4.

8. Select **OK** when completed to update the radio QoS settings for this policy. Select **Reset** to revert the WMM screen back to its last saved configuration.
9. Select the **Admission Control** tab to configure an admission control configuration for selected radio QoS policy. Admission control requires clients send their *traffic specifications* (TSPEC) to a controller managed Access Point before they can transmit or receive data within the controller managed network.

The name of the Radio QoS policy for which the admission control settings apply displays in the banner of the QoS Policy screen.

FIGURE 205 Radio QoS Policy Admission Control screen

10. Select the **Enable admission control for firewall Detected Traffic (e.g., SIP)** check box to apply Radio QoS settings to traffic detected by the controller Firewall. This feature is enabled by default.
11. Select the **Implicit TPSEC** check box to require wireless clients to send their traffic specifications to a controller managed Access Point before they can transmit or receive data. If enabled, this setting applies to just this radio's QoS policy. This feature is enabled by default.
12. Set the following **Voice Access** admission control settings for this radio QoS policy:

Enable Voice

Select the check box to enable admission control for this policy's voice traffic. Only voice traffic admission control is enabled, not any of the other access categories (each access category must be separately enabled and configured). The default is 75.

Maximum Airtime

Set the maximum airtime (in the form of a percentage of the radio's bandwidth) allotted to admission control for voice supported client traffic. The available percentage range is from 0-150%, with 150% being available to account for over-subscription. This value helps ensure the radio's bandwidth is available for high bandwidth voice traffic (if anticipated on the wireless medium) or other access category traffic if voice support is not prioritized. Voice traffic requires longer radio airtime to process, so set a longer airtime value if this radio QoS policy is intended to support voice. The default value is 75.

Maximum Wireless Clients	Set the number of voice supported wireless clients allowed to exist (and consume bandwidth) within the radio's QoS policy. Select from an available range of 0-256 clients. Consider setting this value proportionally to the number of other QoS policies supporting the voice access category, as wireless clients supporting voice use a greater proportion of controller resources than lower bandwidth traffic (like low and best effort categories). The default value is 100 clients.
Maximum Roamed Wireless Clients	Set the number of voice supported wireless clients allowed to roam to a different controller managed access point radio. Select from a range of 0-256 clients. The default value is 10 roamed clients.
Reserved for Roam	Set the roam utilization (in the form of a percentage of the radio's bandwidth) allotted to admission control for voice supported clients who have roamed to a different controller managed access point radio. The available percentage range is from 0-150%, with 150% available to account for over-subscription. The default value is 10.

13. Set the following **Normal (Best Effort) Access** admission control settings for this radio QoS policy

Enable Best Effort	Select the check box to enable admission control for this policy's video traffic. Only normal background traffic admission control is enabled, not any of the other access categories (each access category must be separately enabled and configured). This feature is disabled by default.
Maximum Airtime	Set the maximum airtime (in the form of a percentage of the radio's bandwidth) allotted to admission control for normal background client traffic. The available percentage range is from 0-150%, with 150% being available to account for over-subscription. This value helps ensure the radio's bandwidth is available for lower bandwidth normal traffic (if anticipated to proliferate the wireless medium). Normal background traffic only needs a short radio airtime to process, so set an intermediate airtime value if this radio QoS policy is reserved for background data support. The default value is 75.
Maximum Wireless Clients	Set the number of wireless clients supporting background traffic allowed to exist (and consume bandwidth) within the radio's QoS policy. Select from an available range of 0-256 clients. The default value is 100 clients.
Maximum Roamed Wireless Clients	Set the number of normal background supported wireless clients allowed to roam to a different controller managed access point radio. Select from a range of 0-256 clients. The default value is 10 roamed clients.
Reserved for Roam	Set the roam utilization (in the form of a percentage of the radio's bandwidth) allotted to admission control for normal background supported clients who have roamed to a different controller managed radio. The available percentage range is from 0-150%, with 150% available to account for over-subscription. The default value is 10%.

14. Set the following **Video Access** admission control settings for this radio QoS policy:

Enable Video	Select the check box to enable admission control for this policy's video traffic. Only video traffic admission control is enabled, not any of the other access categories (each access category must be separately enabled and configured). This feature is disabled by default.
Maximum Airtime	Set the maximum airtime (in the form of a percentage of the radio's bandwidth) allotted to admission control for video supported client traffic. The available percentage range is from 0-150%, with 150% being available to account for over-subscription. This value helps ensure the radio's bandwidth is available for high bandwidth video traffic (if anticipated on the wireless medium) or other access category traffic if video support is not prioritized. Video traffic requires longer radio airtime to process, so set a longer airtime value if this radio QoS policy is intended to support video. The default value is 75.
Maximum Wireless Clients	Set the number of video supported wireless clients allowed to exist (and consume bandwidth) within the radio's QoS policy. Select from an available range of 0-256 clients. Consider setting this value proportionally to the number of other QoS policies supporting the video access category, as wireless clients supporting video use a greater proportion of controller resources than lower bandwidth traffic (like low and best effort categories).
Maximum Roamed Wireless Clients	Set the number of video supported wireless clients allowed to roam to a different controller managed access point radio. Select from a range of 0-256 clients. The default value is 10 roamed clients.
Reserved for Roam	Set the roam utilization (in the form of a percentage of the radio's bandwidth) allotted to admission control for video supported clients who have roamed to a different controller managed radio. The available percentage range is from 0-150%, with 150% accounting for over-subscription. The default value is 10.

15. Set the following **Low (Background) Access** admission control settings for this radio QoS policy

Enable Background	Select the check box to enable admission control for this policy's lower priority best effort traffic. Only low best effort traffic admission control is enabled, not any of the other access categories (each access category must be separately enabled and configured).
Maximum Airtime	Set the maximum airtime (in the form of a percentage of the radio's bandwidth) allotted to admission control for low, best effort, client traffic. The available percentage range is from 0-150%, with 150% being available to account for over-subscription. Best effort traffic only needs a short radio airtime to process, so set an intermediate airtime value if this radio QoS policy is reserved to support background data. The default value is 75%.
Maximum Wireless Clients	Set the number of low and best effort supported wireless clients allowed to exist (and consume bandwidth) within the radio's QoS policy. Select from an available range of 0-256 clients. The default value is 100 clients.
Maximum Roamed Wireless Clients	Set the number of low and best effort supported wireless clients allowed to roam to a different controller managed access point radio. Select from a range of 0-256 clients. The default value is 10 roamed clients.
Reserved for Roam	Set the roam utilization (in the form of a percentage of the radio's bandwidth) allotted to admission control for normal background supported clients who have roamed to a different controller managed access point radio. The available percentage range is from 0-150%, with 150% available to account for over-subscription. The default value is 10%.

16. Select the **Multimedia Optimization** tab to configure advanced multimedia QoS configuration for selected radio QoS policy.

The screenshot shows the 'Multimedia Optimizations' tab in a configuration window. Under the 'Accelerated Multicast' section, there are five settings, each with a help icon (i), a value field, and a range in parentheses:

- Maximum number of wireless clients allowed: 25 (0 to 256)
- When wireless client count exceeds the above limit: reject
- Maximum multicast streams per client: 2 (1 to 4)
- Packets per second for multicast flow for it to be accelerated: 25 (1 to 500)
- Timeout for wireless clients: 60 (5 to 6,000)

At the bottom right, there are three buttons: 'OK', 'Reset', and 'Exit'.

FIGURE 206 Radio QoS Policy Multimedia Optimizations screen

17. Set the following **Accelerated Multicast** settings for this radio QoS policy:

Maximum number of wireless clients allowed

Specify the maximum number of wireless clients (between 0 and 256) allowed to use accelerated multicast. The default value is 25.

When wireless client count exceeds the above limit

When the wireless client count using accelerated multicast exceeds the maximum number set the radio to either reject new wireless clients or to revert existing clients to a non-accelerated state.

Maximum multicast streams per client

Specify the maximum number of multicast streams (between 1 and 4) wireless clients can use. The default value is 2.

Packets per second for multicast flow for it to be accelerated

Specify the threshold of multicast packets per second (between 1 and 500) that triggers acceleration for wireless clients. The default value is 25.

Timeout for wireless clients

Specify a timeout value in seconds (between 5 and 6,000) for wireless clients to revert back to a non-accelerated state. The default value is 60.

18. Select **OK** to update the radio QoS admission control settings for this policy. Select **Reset** to revert to the last saved configuration.

Radio QoS Configuration and Deployment Considerations

[Radio QoS Policy](#)

Before defining a radio QoS policy, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- To support QoS, each multimedia application, wireless client and WLAN is required to support WMM.
- WMM enabled clients can co-exist with non-WMM clients on the same WLAN. Non-WMM clients are always assigned a Best Effort access category.
- Brocade Mobility recommends default WMM values be used for all deployments. Changing these values can lead to unexpected traffic blockages, and the blockages might be difficult to diagnose.
- Overloading an Access Point radio with too much high priority traffic (especially voice) degrades overall controller service quality for all users.
- TSPEC admission control is only available with newer voice over WLAN phones. Many legacy voice devices do not support TPSEC or even support WMM traffic prioritization.

AAA Policy

Authentication, Authorization, and Accounting (AAA) provides the mechanism network administrators define access control within the controller managed network.

The controller can interoperate with external RADIUS and LDAP Servers (AAA Servers) to provide user database information and user authentication data. Each WLAN controller managed by the controller can maintain its own unique AAA configuration.

AAA provides a modular way of performing the following services:

Authentication – Authentication provides a means for identifying users, including login and password dialog, challenge and response, messaging support and (depending on the security protocol), encryption. Authentication is the technique by which a user is identified before allowed access to the controller managed network. Configure AAA authentication by defining a list of authentication methods, and then applying the list to various interfaces. The list defines the authentication schemes performed and their sequence. The list must be applied to an interface before the defined authentication technique is conducted.

Authorization – Authorization occurs immediately after authentication. Authorization is a method for remote access control, including authorization for services and individual user accounts and profiles. Authorization functions through the assembly of attribute sets describing what the user is authorized to perform. These attributes are compared to information contained in a database for a given user and the result is returned to AAA to determine the user's actual capabilities and restrictions. The database could be located locally on the controller or be hosted remotely on a RADIUS server. Remote RADIUS servers authorize users by associating *attribute-value* (AV) pairs with the appropriate user. Each authorization method must be defined through AAA. When AAA authorization is enabled it's applied equally to all interfaces on the controller managed network.

Accounting – Accounting is the method for collecting and sending security server information for billing, auditing, and reporting user data; such as start and stop times, executed commands (such as PPP), number of packets, and number of bytes. Accounting enables wireless network administrators to track the services users are accessing and the network resources they are consuming. When accounting is enabled, the network access server reports user activity to a RADIUS security server in the form of accounting records. Each accounting record is comprised of

AV pairs and is stored on the controller’s access control server. The data can be analyzed for network management, client billing, and/or auditing. Accounting methods must be defined through AAA. When AAA accounting is activated for the controller, it’s applied equally to all interfaces on the controller’s access servers.

To define unique controller WLAN AAA configurations:

1. Select **Configuration > Wireless > AAA Policy** to display existing AAA policies.

The **Authentication, Authorization, and Accounting (AAA)** screen lists those AAA policies created thus far. Any of these policies can be selected and applied to the controller.

AAA Policy	Accounting Packet Type	Request Interval	IAC Policy	Server Pooling Mode
policy1	Start/Stop	30m 0s		Fallover

Type to search in tables Row Count: 1

Add **Edit** **Delete**

FIGURE 207 Authentication, Authorization, and Accounting (AAA) screen

2. Refer to the following information listed for each existingAAA policy:

- AAA Policy** Displays the name assigned to the AAA policy when it was initially created. The name cannot be edited within a listed profile.
- Accounting Packet Type** Displays the accounting type set for the AAA policy. Options include:
 - Start Only* - Sends a start accounting notice to initiate user accounting.
 - Start/Stop* - Sends a start accounting notice at the beginning of a process and a stop notice at the end of a process. The start accounting record is sent in the background. The requested process begins regardless of whether the start accounting notice is received by the accounting server.

Request Timeout	Displays the time between 1 and 60 seconds for the wireless controller's re-transmission of request packets. The default is 3 seconds. If this time is exceeded, the authentication session is terminated.
DSCP	Displays the DSCP value as a 6-bit parameter in the header of every IP packet used for packet classification. The valid range is between 0 and 63 with a default value of 46.
NAI Routing Enable	Displays NAI routing status. AAA servers identify clients using the NAI. The NAI is a character string in the format of an e-mail address as either user or user@ but it need not be a valid e-mail address or a fully qualified domain name. The NAI can be used either in a specific or generic form. The specific form, which must contain the user portion and may contain the @ portion, identifies a single user. The generic form allows all users in a given or without a to be configured on a single command line. Each user still needs a unique security association, but these associations can be stored on a AAA server. The original purpose of the NAI was to support roaming between dialup ISPs. Using NAI, each ISP need not have all the accounts for all of its roaming partners in a single RADIUS database. RADIUS servers can proxy requests to remote servers for each.
NAC Enable	A green checkmark defines NAC as enabled, while a Red X defines NAC disabled with this AAA policy.

5. Select an item from the table and click **Edit** or click **Add** to create a new policy.

The screenshot shows a configuration window titled "Server Id 1". It contains the following settings:

- Host:** 192.158.91.6 (IP Address)
- Port:** 1812 (1 to 65,535)
- Server Type:** Host
- Secret:** [Masked] (Show checkbox)
- Request Proxy Mode:** None
- Request Attempts:** 3 (1 to 10)
- Request Timeout:** 3 (Seconds) (1 to 60)
- Retry Timeout Factor:** 100 (50 to 200)
- DSCP:** 46 (0 to 63)
- Network Access Identifier Routing:**
 - NAI Routing Enable:
 - Realm: [Empty text box]
 - Realm Type: Prefix Suffix
 - Strip Realm:

Buttons at the bottom: OK, Reset, Exit.

FIGURE 209 AAA Policy - Add RADIUS Authentication Server

6. Define the following settings to add or modify new AAA RADIUS authentication server configuration

Server ID	Define the numerical server index (1-6) for the authentication server when added to the list available to the wireless controller.
Host	Specify the IP address or hostname of the RADIUS authentication server.
Port	Define or edit the port on which the RADIUS server listens to traffic within the wireless controller managed network. The port range is 1 to 65,535. The default port is 1.
Server Type	Select the type of AAA server in use either Host, onboard-self, or onboard-controller.

Secret	Specify the secret used for authentication on the selected RADIUS server. By default the secret will be displayed as asterisks. To show the secret in plain text, check the Show box.
Request Proxy Mode	Select the method of proxy that browsers communicate with the RADIUS authentication server. The mode could either be None, Through Wireless Controller, or Through RF Domain.
Request Attempts	Specify the number of attempts a client can retransmit a missed frame to the RADIUS server before it times out of the authentication session. The available range is between 1 and 10 attempts. The default is 3 attempts.
Request Timeout	Specify the time between 1 and 60 seconds for the wireless controller's re-transmission of request packets. If this time is exceeded, the authentication session is terminated.
Retry Timeout Factor	Specify the amount of time between 50 and 200 seconds between retry timeouts for the wireless controller's re-transmission of request packets. The default is 100.
DSCP	Specify the DSCP value as a 6-bit parameter in the header of every IP packet used for packet classification. The valid range is between 0 and 63 with a default value of 46.
NAI Routing Enable	Check to enable NAI routing. AAA servers identify clients using the NAI. The NAI is a character string in the format of an e-mail address as either user or user@ but it need not be a valid e-mail address or a fully qualified domain name. The NAI can be used either in a specific or generic form. The specific form, which must contain the user portion and may contain the @ portion, identifies a single user. The generic form allows all users in a given or without a to be configured on a single command line. Each user still needs a unique security association, but these associations can be stored on a AAA server. The original purpose of the NAI was to support roaming between dialup ISPs. Using NAI, each ISP need not have all the accounts for all of its roaming partners in a single RADIUS database. RADIUS servers can proxy requests to remote servers for each.
Realm	Enter the realm name in the field. The name cannot exceed 50 characters. When the controller's RADIUS server receives a request for a user name the server references a table of usernames. If the user name is known, the server proxies the request to the RADIUS server.
Realm Type	Specify the type of realm that is being used, either a Prefix or a Suffix.
Strip Realm	Check strip to remove information from the packet when NAI routing is enabled.

7. Select the **RADIUS Accounting** tab and refer to the following information about configured RADIUS Accounting profiles.

RADIUS Authentication RADIUS Accounting Settings								
Server Id	Host	Port	Server Type	Request Timeout	Request Attempts	DSCP	Request Proxy Mode	NAI Routing Enable
1	158.89.84.189	1	Host	5s	3	34	None	<input checked="" type="checkbox"/>

FIGURE 210 AAA Policy - RADIUS Accounting screen

Server ID	Displays the numerical server index (1-6) for the accounting server when added to the list available to the wireless controller.
Host	Displays the IP address or hostname of the RADIUS authentication server.
Port	Displays the port on which the RADIUS server listens to traffic within the wireless controller managed network. The port range is 1 to 65,535. The default port is 1.
Server Type	Displays the type of AAA server in use either Host, onboard-self, or onboard-controller.
Request Timeout	Displays the time between 1 and 60 seconds for the wireless controller's re-transmission of request packets. If this time is exceeded, the authentication session is terminated.
Request Attempts	Displays the number of attempts a client can retransmit a missed frame to the RADIUS server before it times out of the authentication session. The available range is between 1 and 10 attempts. The default is 3 attempts.

DSCP	Displays the DSCP value as a 6-bit parameter in the header of every IP packet used for packet classification. The valid range is between 0 and 63 with a default value of 46.
Request Proxy Mode	Displays the method of proxy that browsers communicate with the RADIUS authentication server. The mode could either be <i>None</i> , <i>Through Wireless Controller</i> , or <i>Through RF Domain Manager</i> .
NAI Routing Enable	Displays NAI routing status. AAA servers identify clients using the NAI. The NAI is a character string in the format of an e-mail address as either <i>user</i> or <i>user@</i> but it need not be a valid e-mail address or a fully qualified domain name. The NAI can be used either in a specific or generic form. The specific form, which must contain the user portion and may contain the @ portion, identifies a single user. The generic form allows all users in a given or without a to be configured on a single command line. Each user still needs a unique security association, but these associations can be stored on a AAA server. The original purpose of the NAI was to support roaming between dialup ISPs. Using NAI, each ISP need not have all the accounts for all of its roaming partners in a single RADIUS database. RADIUS servers can proxy requests to remote servers for each.

- To edit an existing accounting profile, select the profile and click **Edit**. To add a new policy click **Add**.

FIGURE 211 AAA Policy - Add RADIUS Accounting Server

Server ID	Displays the numerical server index (1-6) for the accounting server when added to the list available to the access point.
Host	Specify the IP address or hostname of the RADIUS authentication server.
Port	Define or edit the port on which the RADIUS server listens to traffic within the controller managed network. The port range is 1 to 65,535. The default port is 1813.
Server Type	Select the type of AAA server as either <i>Host</i> , <i>onboard-self</i> , or <i>onboard-controller</i> .
Secret	Specify the secret (password) used for authentication on the selected RADIUS server. By default the secret is displayed as asterisks.

Request Proxy Mode	Select the method of proxy that browsers communicate with the RADIUS authentication server. The mode could either be <i>None</i> , <i>Through Wireless Controller</i> , or <i>Through RF Domain Manager</i> .
Request Attempts	Displays the number of attempts a client can retransmit a missed frame to the RADIUS server before it times out of the authentication session. The available range is between 1 and 10 attempts. The default is 3 attempts.
Request Timeout	Specify the time for the access point's re-transmission of request packets. The default is 5 seconds. If this time is exceeded, the authentication session is terminated.
Retry Timeout Factor	Specify the amount of time between 50 and 200 seconds between retry timeouts for the access points re-transmission of request packets. The default is 100.
DSCP	Displays the DSCP value as a 6-bit parameter in the header of every IP packet used for packet classification. The valid range is between 0 and 63 with a default value of 34.
NAI Routing Enable	Displays NAI routing status. AAA servers identify clients using the NAI. The NAI is a character string in the format of an e-mail address as either user or user@ but it need not be a valid e-mail address or a fully qualified domain name. The NAI can be used either in a specific or generic form. The specific form, which must contain the user portion and may contain the @ portion, identifies a single user. The generic form allows all users in a given or without a to be configured on a single command line. Each user still needs a unique security association, but these associations can be stored on a AAA server. The original purpose of the NAI was to support roaming between dialup ISPs. Using NAI, each ISP need not have all the accounts for all of its roaming partners in a single RADIUS database. RADIUS servers can proxy requests to remote servers for each.
Realm	Enter the realm name. The name cannot exceed 64 characters. When the access point's RADIUS server receives a request for a user name the server references a table of usernames. If the user name is known, the server proxies the request to the RADIUS server.
Realm Type	Specify the realm as either <i>Prefix</i> or <i>Suffix</i> .
Strip Realm	Select the radio button to remove information from the packet when NAI routing is enabled.

9. Click the **Settings** tab and configure to the following information:

The screenshot shows the 'AAA Policy - Settings' screen with the following configurations:

- RADIUS Authentication:** Protocol for MAC, Captive-Portal Authentication is set to PAP.
- RADIUS Accounting:** Accounting Packet Type is Start/Stop; Request Interval is 30 Minutes; Accounting Server Preference is Prefer Same Authentication Server Host.
- RADIUS Address Format:** Format is Dash Delimiter (aa-bb-cc-dd-ee-ff); Case is Uppercase; Attributes is Username / Password.
- Server Pooling:** Server Pooling Mode is Failover.
- EAP Wireless Client Settings:** Client Attempts is 3; Request Timeout is 3 seconds; ID Request Timeout is 10 seconds; Retransmission Scale Factor is 100.

FIGURE 212 AAA Policy - Settings screen

Protocol for MAC, Captive-Portal Authentication	The authentication protocol <i>Password Authentication Protocol (PAP)</i> or <i>Challenge Handshake Authentication Protocol (CHAP)</i> when the server is used for any non-EAP authentication. PAP is the default setting
Accounting Packet Type	Set the type of RADIUS Accounting Request packets generated. Options include <i>Stop Only</i> , <i>Start/Stop</i> , <i>Start/Interim/Stop</i> . <i>Start/Stop</i> is the default setting
Request Interval	Set the periodicity of the interim accounting requests. The default is 30 minutes.
Accounting Server Preference	Select the server preference for RADIUS Accounting. The options are: <i>Prefer Same Authentication Server Host</i> - Uses the authentication server host name as the host used for RADIUS accounting. This is the default setting. <i>Prefer Same Authentication Server Index</i> - Uses the same index as the authentication server for RADIUS accounting. <i>Select Accounting Server Independently</i> - Allows users to specify a RADIUS accounting server separate from the RADIUS authentication server.
Format	Select the format of the MAC address used in the RADIUS accounting packets.
Case	Lists whether the MAC address is sent using <i>uppercase</i> or <i>lowercase</i> letters. The default setting is uppercase
Attributes	Lists whether the format specified applies only to the username/password in mac-auth or for all attributes that include a MAC address, such as calling-station-id or called-station-id.
Server Pooling Mode	Controls how requests are transmitted across RADIUS servers. <i>Failover</i> implies traversing the list of servers if any server is unresponsive. <i>Load Balanced</i> means using all servers in a round-robin fashion. The default setting is Failover.
Client Attempts	Defines the number of times (1 - 10) an EAP request is transmitted to a Wireless Client before giving up. The default setting is 3.

Request Timeout	Defines the amount of time after which an EAP Request to a Wireless Client is retried. The default setting is 3 seconds.
ID Request Timeout	Defines the amount of time (1 - 60 seconds) after which an EAP ID Request to a Wireless Client is retried. The default setting is 10 seconds.
Retransmission Scale Factor	Configures the scaling of the retransmission attempts. Timeout at each attempt is a function of the Request Timeout factor and Client Attempts number. 100 (default setting) implies a constant timeout at each retry; smaller values indicate more aggressive (shorter) timeouts, larger numbers indicate more conservative (longer) timeouts on each successive attempt.

Association ACL

An Association ACL is a policy-based ACL that either prevents or allows wireless clients from connecting to a controller managed WLAN.

An Association ACL affords a system administrator the ability to grant or restrict client access by specifying a wireless client MAC address or range of MAC addresses to either include or exclude from controller connectivity.

Association ACLs are applied to WLANs as an additional access control mechanism. They can be applied to WLANs from within a WLAN Policy's Advanced configuration screen. For more information on applying an existing Association ACL to a WLAN, see *Configuring Advanced WLAN Settings on page 6-261*.

To define an Association ACL deployable with a controller WLAN:

1. Select **Configuration > Wireless > Association ACL** to display existing Association ACLs.

The **Association Access Control List (ACL)** screen lists those Association ACL policies created thus far. Any of these policies can be selected and applied to the controller.




FIGURE 213 Association Access Control List (ACL) screen

2. Select **Add** to define a new ACL configuration, **Edit** to modify an existing ACL configuration or **Delete** to remove one.

A unique Association ACL screen displays for defining the new ACL or modifying a selected ACL.

Precedence	Starting MAC Address	Ending MAC Address	Allow/Deny
3	00-00-00-00-00-00	FF-FF-FF-FF-FF-FF	Deny

FIGURE 214 Association Access Control List (ACL) screen

3. Select the **+ Add Row** button to add an association ACL template.
4. Set the following parameters for the creation or modification of the Association ACL:

Association ACL	If creating an new Association ACL, provide a name specific to its function. Avoid naming it after a WLAN it may support. The name cannot exceed 32 characters.
Precedence	The rules within a WLAN's ACL are applied to packets based on their precedence values. Every rule has a unique sequential precedence value you define. You cannot add two rules's with the same precedence. The default precedence is 1, so be careful to prioritize ACLs accordingly as they are added.
Starting MAC Address	Provide a starting MAC address for clients requesting association.
Ending MAC Address	Provide an ending MAC address for clients requesting association.
Allow/Deny	Use the drop-down menu to either <i>Allow</i> or <i>Deny</i> access if a MAC address matches this rule.

5. Select the **+ Add Row** button to add MAC address ranges and allow/deny designations.
6. Select **OK** to update the Association ACL settings. Select **Reset** to revert to the last saved configuration.

Association ACL Deployment Considerations

Association ACL

Before defining an Association ACL configuration and applying it to a WLAN, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- Brocade Mobility recommends using the Association ACL screen strategically to name and configure ACL policies meeting the requirements of the particular WLANs they may map to. However, be careful not to name ACLs after specific WLANs, as individual ACL policies can be used by more than one WLAN.
- You cannot apply more than one MAC based ACL to a Layer 2 interface. If a MAC ACL is already configured on a Layer 2 interface, and a new MAC ACL is applied to the interface, the new ACL replaces the previously configured one.

Smart RF Policy

Self Monitoring At Run Time RF Management (Smart RF) is a Brocade Mobility innovation designed to simplify RF configurations for new deployments, while (over time) providing on-going deployment optimization radio performance improvements.

A controller Smart RF policy can reduce deployment costs by scanning the RF environment to determine the best channel and transmit power configuration for each wireless controller managed radio. Smart RF policies can be applied to specific RF Domains, to apply site specific deployment configurations and self-healing values to groups of devices within pre-defined physical RF coverage areas.

Smart RF centralizes the decision process and makes intelligent RF configuration decisions using information obtained from the RF environment. Smart RF helps reduce ongoing management and maintenance costs through the periodic re-calibration of the managed network. Re-calibration can be initiated manually or can be automatically scheduled to ensure the RF configuration is optimized to factor for RF environment changes (such as new sources of interference, or neighboring APs).

Smart RF also provides self-healing functions by monitoring the managed network in real-time and provides automatic mitigation from potentially problematic events such as radio interference, coverage holes and radio failures. Smart RF employs self-healing to enable a WLAN to better maintain wireless client performance and site coverage during dynamic RF environment changes, which typically require manual reconfiguration to resolve.

Smart RF is supported on any RF Domain manager. In standalone environments, the individual controller manages the calibration and monitoring phases. In clustered environments, a single controller is elected a Smart RF master and the remaining cluster members operate as Smart RF clients. In cluster operation, the Smart RF master co-ordinates the calibration and configuration and during the monitoring phase receives information from the Smart RF clients.

NOTE

RF planning must be performed to ensure overlapping coverage exists at a deployment site for Smart RF to be a viable network performance tool. Smart RF can only provide recovery when Access Points are deployed appropriately. Smart RF is not a solution, it's a temporary measure. Administrators need to determine the root cause of RF deterioration and fix it. Smart RF history/events can assist.

To define a Smart RF policy:

1. Select **Configuration > Wireless > Smart RF Policy** to display existing Smart RF policies.

The Smart RF screen lists those Smart RF policies created thus far. Any of these policies can be selected and applied to the controller.

The user has the option of displaying the configurations of each Smart RF Policy defined thus far, or referring to the **Smart RF Browser** and either selecting individual Smart RF polices or selecting existing RF Domains to review which Smart RF policies have been applied. For more information on how RF Domains function, and how to apply a Smart RF policy, see [About RF Domains on page 8-407](#) and [Managing RF Domains on page 8-408](#).

SMART RF Policy	SMART RF Policy Enable	Interference Recovery	Coverage Hole Recovery	Neighbor Recovery	Smart Monitoring Enable
smartr1	X	✓	✓	✓	✓
smartr2	X	✓	✓	✓	✓
smartr3	X	✓	✓	✓	✓
smartr4	X	✓	✓	✓	✓

FIGURE 215 Smart RF Policy screen

2. Refer to the following configuration data for existing Smart RF policies:

Smart RF Policy	Displays the name assigned to the Smart RF policy when it was initially created. The name cannot be modified as part of the edit process.
Smart RF Policy Enable	Displays a green check mark if Smart RF has been enabled for the listed policy. A red “X” designates the policy as being disabled.
Interference Recovery	Displays a green check mark if interference recovery has been enabled for the listed policy. A red “X” designates interference recovery being disabled.

Coverage Hole Recovery	Displays a green check mark if coverage hole recovery has been enabled for the listed policy. A red "X" designates coverage hole recovery being disabled.
Neighbor Recovery	Displays a green check mark if neighbor recovery has been enabled for the listed policy. A red "X" designates neighbor recovery being disabled.
Smart Monitoring Enable	Displays a green check mark if smart monitoring has been enabled for the listed policy. A red "X" designates smart monitoring being disabled.

3. Select **Add** to create a new Smart RF policy, **Edit** to modify the attributes of a existing policy or **Delete** to remove obsolete policies from the list of those available.

The **Basic Configuration** screen displays by default for the new or modified Smart RF policy.

Basic Settings

Sensitivity Low Medium High Custom

SMART RF Policy Enable

Auto Assign Sensor

Interference Recovery

Coverage Hole Recovery

Neighbor Recovery

Calibration Assignment

Enable Per Area

Enable Per Floor

OK Reset Exit

FIGURE 216 Smart RF Basic Configuration screen

4. Refer to the **Basic Settings** field to enable a Smart RF policy and define its sensitivity and detector status.

Sensitivity	Select a radio button corresponding to the desired Smart RF sensitivity. Options include <i>Low</i> , <i>Medium</i> , <i>High</i> and <i>Custom</i> . Medium, is the default setting. Select the Custom sensitivity option to enable the Interference Recovery, Coverage Hole Recovery and Neighbor Recovery options as additional Smart RF recovery options.
SMART RF Policy Enable	Select the Smart RF Policy Enable check box to enable this Smart RF policy for immediate controller network support or inclusion with a RF Domain. Smart RF is enabled by default.
Assign Auto Sensor	Select the check box to enable Smart RF to select a radio as a sensor. Sensor radios provide dedicated 24 x 7 dual-band 802.11a/b/g/n monitoring on a single radio to perform real-time monitoring and self healing as needed. Sensor radios are automatically defined and configured during calibration to provide real-time monitoring for self-healing. Sensor radios are dedicated for monitoring only and do not provide client services. This feature is disabled by default.
Interference Recovery	Select the check box to enable Interference Recovery from neighboring radios and other sources of WiFi and non-WiFi interference when excess noise and interference is detected within the Smart RF supported radio coverage area. Smart RF provides mitigation from interference sources by monitoring the noise levels and other RF parameters on an Access Point radio's current channel. When a noise threshold is exceeded, Smart RF can select an alternative channel with less interference. To avoid channel flapping, a hold timer is defined which disables interference avoidance for a specific period of time upon detection. Interference Recovery is enabled by default.
Coverage Hole Recovery	Select the check box to enable Coverage Hole Recovery when a radio coverage hole is detected within the Smart RF supported radio coverage area. When coverage hole is detected, Smart RF first determines the power increase needed based on the signal to noise ratio for a client as seen by the Access Point radio. If a client's signal to noise value is above the threshold, the transmit power is increased until the signal to noise rate falls below the threshold.
Neighbor Recovery	Select the check box to enable Neighbor Recovery when a failed radio is detected within the Smart RF supported radio coverage area. Smart RF can provide automatic recovery by instructing neighboring APs to increase their transmit power to compensate for the coverage loss. Neighbor recovery is enabled by default when the sensitivity setting is medium.

5. Refer to the **Calibration Assignment** field to define whether Smart RF Calibration and radio grouping is conducted by area or building. Both options are disabled by default.
6. Select **OK** to update the Smart RF Basic Configuration settings for this policy. Select **Reset** to revert to the last saved configuration.
7. Select **Channel and Power**.

Use the Channel and Power screen to refine Smart RF power settings over both 5 and 2.4 GHz radios and select channel settings in respect to the device channel usage.

Power Settings

5 GHz Minimum Power (1 to 20 dBm)

5 GHz Maximum Power (1 to 20 dBm)

2.4 GHz Minimum Power (1 to 20 dBm)

2.4 GHz Maximum Power (1 to 20 dBm)

Channel Settings

5 GHz Channels Select

5 GHz Channel Width 20MHz 40MHz Automatic

2.4 GHz Channels Select

2.4 GHz Channel Width 20MHz 40MHz Automatic

OK Reset Exit

FIGURE 217 Smart RF Channel and Power screen

NOTE

The Power Settings and Channel Settings parameters are only enabled when Custom or Medium is selected as the Sensitivity setting from the Basic Configuration screen.

8. Refer to the **Power Settings** field to define Smart RF recovery settings for either the selected 5.0 GHz (802.11a) or 2.4 GHz (802.11bg) radio.

5.0 GHz Minimum Power

Use the spinner control to select a 1 - 20 dBm minimum power level for Smart RF to assign to a radio in the 5 GHz band. 4 dBm is the default setting.

5.0 GHz Maximum Power

Use the spinner control to select a 1 - 20 dBm maximum power level Smart RF can assign a radio in the 5 GHz band. 17 dBm is the default setting.

2.4 GHz Minimum Power

Use the spinner control to select a 1 - 20 dBm minimum power level Smart RF can assign a radio in the 2.4 GHz band. 4 dBm is the default setting.

2.4 GHz Maximum Power

Use the spinner control to select a 1 - 20 dBm maximum power level Smart RF can assign a radio in the 2.4 GHz band. 17 dBm is the default setting.

9. Set the following **Channel Settings** for the 5.0 GHz and 2.4 GHz radios:

5.0 GHz Channels

Use the **Select** drop-down menu to select the 5 GHz channels used in Smart RF scans.

5.0 Channel Width

20 and 40 MHz channel widths are supported by the 802.11a radio. 20/40 MHz operation (the default setting for the 5 GHz radio) allows the Access Point to receive packets from clients using 20 MHz of bandwidth while transmitting a packet using 40 MHz bandwidth. This mode is supported for 11n users on both the 2.4 and 5 GHz radios. If an 11n user selects two channels (a Primary and Secondary channel), the system is configured for dynamic 20/40 operation. When 20/40 is selected, clients can take advantage of *wider channels*. 802.11n clients experience improved throughput using 40 MHz while legacy clients (either 802.11a or 802.11b/g depending on the radio selected) can still be serviced without interruption using 20 MHz. Select **Automatic** to enable automatic assignment of channels to working radios to avoid channel overlap and avoid interference from external RF sources. 40MHz is the default setting.

2.4 GHz Channels

Set the 2.4 GHz channels used in Smart RF scans.

2.4 GHz Channel Width

20 and 40 MHz channel widths are supported by the 802.11a radio. 20 MHz is the default setting for 2.4 GHz radios. 20/40 MHz operation (the default setting for the 5 GHz radio) allows the Access Point to receive packets from clients using 20 MHz of bandwidth while transmitting a packet using 40 MHz bandwidth. This mode is supported for 11n users on both the 2.4 and 5 GHz radios. If an 11n user selects two channels (a Primary and Secondary channel), the system is configured for dynamic 20/40 operation. When 20/40 is selected, clients can take advantage of *wider channels*. 802.11n clients experience improved throughput using 40 MHz while legacy clients (either 802.11a or 802.11b/g depending on the radio selected) can still be serviced without interruption using 20 MHz. Select **Automatic** to enable automatic assignment of channels to working radios to avoid channel overlap and avoid interference from external RF sources. 20MHz is the default setting.

10. Select **OK** to update the Smart RF Channel and Power settings for this policy. Select **Reset** to revert to the last saved configuration.

11. Select the **Scanning Configuration** tab.

Monitoring Configuration

Smart Monitoring Enable

Scanning Configuration for 5.0 GHz

Duration: 50 (20 to 150 milliseconds)

Frequency: 6 Seconds (1 to 120)

Extended Scan Frequency: 5 (0 to 50)

Sample Count: 5 (1 to 15)

Client Aware Scanning: 0 (0 to 255)

Power Save Aware Scanning: Dynamic Strict Disable

Voice Aware Scanning: Dynamic Strict Disable

Scanning Configuration for 2.4 GHz

Duration: 50 (20 to 150 milliseconds)

Frequency: 6 Seconds (1 to 120)

Extended Scan Frequency: 5 (0 to 50)

Sample Count: 5 (1 to 15)

Client Aware Scanning: 0 (0 to 255)

Power Save Aware Scanning: Dynamic Strict Disable

Voice Aware Scanning: Dynamic Strict Disable

Note: The system automatically configures optimum values for certain fields, if you select the sensitivity option under 'Basic Settings' as 'Low', 'Medium' or 'High'. Some of the SMART RF parameters appear disabled in this case. Please choose the 'Custom' sensitivity option to enable the fields and manually enter each value.

OK Reset Exit

FIGURE 218 Smart RF Scanning Configuration screen

NOTE

The monitoring and scanning parameters within the Scanning Configuration screen are only enabled when Custom is selected as the Sensitivity setting from the Basic Configuration screen.

12. Enable or disable **Smart Monitoring Enable** by selecting the check box. The feature is enabled by default.

When enabled, detector radios monitor their coverage areas for potential failed peers or coverage area holes requiring transmission adjustments for coverage compensation.

13. Set the following **Scanning Configurations** for both the 2.4 and 5 GHz radio bands:

Duration	Set a channel scan duration (between 20 - 150 milliseconds) Access Point radios use to monitor devices within the managed network and, if necessary, perform self healing and neighbor recovery to compensate for coverage area losses within a RF Domain. The default setting is 50 milliseconds for both the 2.4 and 5 GHz bands.
Frequency	Set the scan frequency using the drop-down menu. Set a scan frequency in either <i>Seconds</i> (1 - 120) or <i>Minutes</i> (0 - 2). The default setting is 6 seconds for both the 5 and 2.4 GHz bands.
Extended Scan Frequency	Use the spinner control to set an extended scan frequency between 0 - 50. This is the frequency radios scan channels on other than their peer radios. The default setting is 5 for both the 5 and 2.4 GHz bands.

Sample Count	Use the spinner control to set a sample scan count value between 1 - 15. This is the number of RF readings radios gather before they send the data to the Smart RF master. The default setting is 5 for both the 5 and 2.4 GHz bands
Power Save Aware Scanning	Select either the <i>Dynamic</i> , <i>Strict</i> or <i>Disable</i> radio button to define how power save scanning is set for Smart RF. Strict disables smart monitoring as long as a power save capable client is associated to a radio. Dynamic disables smart monitoring as long as there is data buffered for a power save client at the radio. The default setting is Dynamic for both the 5 and 2.4 GHz bands.
Voice Aware	Select either the <i>Dynamic</i> , <i>Strict</i> or <i>Disable</i> radio button to define how voice aware recognition is set for Smart RF. Strict disables smart monitoring as long as a voice client is associated to a radio. Dynamic disables smart monitoring as long as there is data buffered for a voice client at the radio. The default setting is Dynamic for both the 5 and 2.4 GHz bands.

14. Select **OK** to update the Smart RF Scanning Configuration settings for this policy. Select **Reset** to revert to the last saved configuration.

15. Select **Advanced Configuration**.

The **Neighbor Recovery** tab displays by default. Use the *Neighbor*, *Interference* and *Coverage Hole* recovery tabs to define how 5 and 2.4 GHz radios compensate for failed neighbor radios, interference impacting the Smart RF support controller network and detected coverage holes requiring neighbor radio intervention.

NOTE

The recovery parameters within the Neighbor Recovery, Interference and Coverage Hole Recovery tabs are only enabled when Custom is selected as the Sensitivity setting from the Basic Configuration screen.

16. Set the following **Neighbor Recovery** variables for the Smart RF configuration:

The screenshot shows the 'Neighbor Recovery' tab of the Smart RF Advanced Configuration screen. It includes sections for 'Hold Time', 'Neighbor Recovery', and 'Dynamic Sample Recovery'. The 'Power Hold Time' is set to 0 seconds. The '5 GHz Neighbor Power Threshold' and '2.4 GHz Neighbor Power Threshold' are both set to -70 dBm. Under 'Dynamic Sample Recovery', 'Dynamic Sample Enabled' is unchecked, 'Dynamic Sample Retries' is set to 3, and 'Dynamic Sample Threshold' is set to 5. A note at the bottom states: 'Note: The system automatically configures optimum values for certain fields, if you select the sensitivity option under 'Basic Settings' as 'Low', 'Medium' or 'High'. Some of the SMART RF parameters appear disabled in this case. Please choose the 'Custom' sensitivity option to enable the fields and manually enter each value.'

FIGURE 219 Smart RF Advanced Configuration screen - Neighbor Recovery tab

Power Hold Time Defines the minimum time between two radio power changes during neighbor recovery. Set the time in either *Seconds* (0 - 3,600), *Minutes* (0 - 60) or *Hours* (0 - 1). The default setting is 0 seconds.

Channel Hold Time Defines the minimum time between channel changes during neighbor recovery. Set the time in either *Seconds* (0 - 86,400), *Minutes* (0 - 1,440) or *Hours* (0 - 24) or *Days* (0 - 1). The default setting is 3,600 seconds.

17. Set the following **Neighbor Recovery** parameters:

5.0 GHz Neighbor Recovery Power Threshold Use the spinner control to set a value between -85 to -55 dBm the 5.0 GHz radio uses as a maximum power increase threshold if the radio is required to increase its output power to compensate for a failed radio within its wireless radio coverage area. The default value is -70 dBm.

2.4 GHz Neighbor Recovery Power Threshold Use the spinner control to set a value between -85 to -55 dBm the 2.4 GHz radio uses as a maximum power increase threshold if the radio is required to increase its output power to compensate for a failed radio within its wireless radio coverage area. The default value is -70 dBm.

18. Select **OK** to update the Smart RF Neighbor Recovery settings for this policy. Select **Reset** to revert to the last saved configuration.

19. Select the **Interference Recovery** tab.

The screenshot shows the 'Interference Recovery' tab in the Smart RF Advanced Configuration screen. The settings are as follows:

- Interference:**
- Noise:**
- Noise Factor:** 1.50 (range: 1.0 - 3.0)
- Channel Hold Time:** 1 Hours (range: 0 to 24)
- Client Threshold:** 50 (range: 1 to 255)
- 5 GHz Channel Switch Delta:** 20 (range: 5 to 35 dBm)
- 2.4 GHz Channel Switch Delta:** 20 (range: 5 to 35 dBm)

Note: The system automatically configures optimum values for certain fields, if you select the sensitivity option under 'Basic Settings' as 'Low', 'Medium' or 'High'. Some of the SMART RF parameters appear disabled in this case. Please choose the 'Custom' sensitivity option to enable the fields and manually enter each value.

Buttons at the bottom: OK, Reset, Exit.

FIGURE 220 Smart RF Advanced Configuration screen - Interference Recovery tab

20. Set the following **Interference Recovery** parameters:

Interference

Select the check box to allow the Smart RF policy to scan for excess interference from supported radio devices. WLANs are susceptible to sources of interference, such as neighboring radios, cordless phones, microwave ovens and Bluetooth devices. When interference for WiFi sources is detected, Smart RF supported devices can change the channel and move to a cleaner channel. This feature is enabled by default.

Noise

Select the check box to allow the Smart RF policy to scan for excess noise from WiFi devices. When detected, Smart RF supported devices can change their channel and move to a cleaner channel. This feature is enabled by default.

Client Threshold

Use the spinner to set a client threshold for the Smart RF policy between 1 - 255. If threshold number of clients are connected to a radio, it does not change its channel even though it requires one, based on the interference recovery determination made by the smart master. The default is 50.

5.0 GHz Channel Switch Delta

Use the spinner to set a channel delta (between 5 - 35 dBm) for the 5.0 GHz radio. This parameter is the difference between noise levels on the current channel and a prospective channel. If the difference is below the configured threshold, the channel will not change. The default setting is 20 dBm.

2.4 GHz Channel Switch Delta

Use the spinner to set a channel delta (between 5 - 35 dBm) for the 2.4 GHz radio. This parameter is the difference between noise levels on the current channel and a prospective channel. If the difference is below the configured threshold, the channel will not change. The default setting is 20 dBm.

21. Select **OK** to update the Smart RF Interference Recovery settings for this policy. Select **Reset** to revert to the last saved configuration.
22. Select the **Coverage Hole Recovery** tab.

The screenshot shows the 'Coverage Hole Recovery' configuration screen. At the top, there are three tabs: 'Neighbor Recovery', 'Interference Recovery', and 'Coverage Hole Recovery'. The 'Coverage Hole Recovery' tab is selected. Below the tabs, there are two sections: 'Coverage Hole Recovery for 5.0 GHz' and 'Coverage Hole Recovery for 2.4 GHz'. Each section contains four parameters: Client Threshold, SNR Threshold, Coverage Interval, and Interval. Each parameter has a spinner control and a range in parentheses. For 5.0 GHz, the values are Client Threshold: 1 (1 to 255), SNR Threshold: 20 (1 to 75 dB), Coverage Interval: 10 (1 to 120), and Interval: 30 (1 to 120). For 2.4 GHz, the values are Client Threshold: 1 (1 to 255), SNR Threshold: 20 (1 to 75 dB), Coverage Interval: 10 (1 to 120), and Interval: 30 (1 to 120). At the bottom of the screen, there are three buttons: 'OK', 'Reset', and 'Exit'. A yellow warning icon and a note are located below the 2.4 GHz section.

Note: The system automatically configures optimum values for certain fields, if you select the sensitivity option under 'Basic Settings' as 'Low', 'Medium' or 'High'. Some of the SMART RF parameters appear disabled in this case. Please choose the 'Custom' sensitivity option to enable the fields and manually enter each value.

FIGURE 221 Smart RF Advanced Configuration screen - Coverage Hole Recovery tab

23. Set the following **Coverage Hole Recovery for 2.4 GHz and 5 GHz** parameters:

Client Threshold	Use the spinner to set a client threshold for the Smart RF policy between 1 - 255. This is the minimum number of clients a radio should have associated in order for coverage hole recovery to trigger. The default setting is 1.
SNR Threshold	Use the spinner control to set a signal to noise threshold (between 1 - 75 dB). This is the signal to noise threshold for an associated client as seen by its associated AP radio. When exceeded, the radio increases its transmit power in order to increase coverage for the associated client. The default value is 20 dB.
Coverage Interval	Define the interval coverage hole recovery should be initiated after a coverage hole is detected. The default is 10 seconds for both the 2.4 and 5.0 GHz radios.
Interval	Define the interval coverage hole recovery should be conducted after a coverage hole is detected. The default is 30 seconds for both the 2.4 and 5.0 GHz radios.

24. Select **OK** to update the Smart RF Coverage Hole Recovery settings for this policy. Select **Reset** to revert to the last saved configuration.

Smart RF Configuration and Deployment Considerations

Smart RF Policy

Before defining a Smart RF policy, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- The Smart RF calibration process impacts associated users and should not be run during business or production hours. The calibration process should be performed during scheduled maintenance intervals or non-business hours.
- For Smart RF to provide effective recovery, RF planning must be performed to ensure overlapping coverage exists at the deployment site. Smart RF can only provide recovery when Access Points are deployed appropriately. Smart RF is not a solution, it's a temporary measure. Administrators need to determine the root cause of RF deterioration and fix it. Smart RF history/events can assist.

Profile Configuration

In this chapter

- [General Profile Configuration](#) 316
- [Profile Cluster Configuration \(Controllers Only\)](#) 318
- [Controller Cluster Profile Configuration and Deployment Considerations](#). 321
- [Profile Adoption Configuration \(APs Only\)](#) 321
- [Profile Interface Configuration](#) 322
- [Profile Network Configuration](#). 350
- [Profile Security Configuration](#) 368
- [Profile Services Configuration](#) 391
- [Profile Management Configuration](#) 393
- [Advanced Profile Configuration](#) 399

Profiles enable administrators to assign a common set of configuration parameters and policies to controllers and Access Points. Profiles can be used to assign common or *unique* network, wireless and security parameters to Wireless Controllers and Access Points across a large, multi segment, site. The configuration parameters within a profile are based on the hardware model the profile was created to support. The controller supports both default and user defined profiles implementing new features or updating existing parameters to groups of Wireless Controllers or Access Points. The central benefit of a profile is its ability to update devices collectively without having to modify individual device configurations.

Profiles assign configuration parameters, applicable policies and WLANs to one or more controllers and Access Points, thus allowing smart administration across large wireless network segments. However, individual devices can still be assigned unique configuration parameters that follow the flat configuration model supported by Brocade Mobility in previous controller software releases. As individual device updates are made, these device no longer share the profile based configuration they originally supported. Changes made to the profile are automatically inherited by all assigned devices, but not those devices who have had their configuration customized. These devices require careful administration, as they no longer can be tracked and as profile members. Their customized configurations overwrite their profile configurations until the profile can be re-applied to the device.

Each controller and Access Point is automatically assigned a default profile unless an AP auto provisioning policy is defined that specifically assigns the Access Point to a user defined profile. A default profile for each supported model is automatically added to a device's configuration file when the device is discovered by the controller. Default profiles can also be manually added prior to discovery when needed. Default profiles are ideal for single site deployments where controllers and Access Points share a common configuration.

Device Model	Default Profile
Brocade Mobility 650 Access Point	default-ap650
AP6511	default-ap6511
AP7131	default-ap7131
RFS4000	default-rfs4000
RFS6000	default-rfs6000
RFS7000	default-rfs7000

User defined profiles are manually created for each supported Wireless Controller and Access Point model. User defined profiles can be manually assigned or automatically assigned to Access Points using an AP Auto provisioning policy. AP Adoption policies provide the means to easily assign profiles to Access Points based on model, serial number, VLAN ID, DHCP option, IP address (subnet) and MAC address.

Brocade Mobility recommends using user defined profiles are useful in larger deployments using centralized Wireless Controllers when groups of devices on different floors, buildings or sites share a common configuration.

Each default and user defined profile contains policies and configuration parameters. Changes made to these parameters are automatically inherited by the devices assigned to the profile.

Review the existing profiles available to the controller to determine whether a new profile requires creation, or an existing profile requires edit or deletion.

To review the existing profiles available to the controller and its supported devices:

1. Select the **Configuration** tab from the Web UI.
2. Select **Profiles** from the Configuration tab.
3. Select **Manage Profiles** from the **Configuration > Profiles** menu.

Advanced WIPS Policy	Lists the name of the Advanced WIPS Policy used with each listed profile to (among other things) block up to 100 client MAC address from controller connectivity.
DHCP Server Policy	Lists the name of the DHCP Server Policy used with each listed profile. The controller's internal DHCP server groups wireless clients based on defined user-class option values. Clients with a defined set of user class values are segregated by class. A DHCP server can associate multiple classes to each pool. Each class in a pool is assigned an exclusive range of IP addresses.
Management Policy	Lists the name of Management policies applied to each listed controller profile. A management policy is a mechanism to allow/deny management access to the controller for separate interfaces and protocols (HTTP, HTTPS, Telnet, SSH or SNMP). Controller management access can be enabled/disabled as required for each policy.
RADIUS Server Policy	Displays the name of the RADIUS Server policy applied to each listed controller profile. A RADIUS Server policy provides customized, profile specific, management of controller authentication data (usernames and passwords).

5. Select the **Add** button to create a new controller profile, **Edit** to revise a selected profile configuration or **Delete** to permanently remove a selected profile.

The following tasks comprise the controller's required profile configuration activities:

- [General Profile Configuration](#)
- [Profile Cluster Configuration \(Controllers Only\)](#)
- [Profile Adoption Configuration \(APs Only\)](#)
- [Profile Interface Configuration](#)
- [Profile Network Configuration](#)
- [Profile Security Configuration](#)
- [Profile Services Configuration](#)
- [Profile Management Configuration](#)
- [Advanced Profile Configuration](#)

General Profile Configuration

Each controller profile requires a provisioning policy and clock synchronization settings as part of its general configuration. Each profile can have a unique provisioning policy and system time.

Wireless Controllers and Access Points are automatically assigned a default profile unless an AP provisioning policy has been defined that specifically assigns Access Points to a user defined profile. During the general configuration process, a provisioning policy can be assigned to a specific profile or a new provisioning policy can be created and applied to the profile. Adoption is the process an AP uses to discover controllers available in the network, pick the most desirable controller, establish an association and obtain its configuration.

Network Time Protocol (NTP) manages time and/or network clock synchronization within the managed network. NTP is a client/server implementation. The Wireless Controller (an NTP client) periodically synchronizes its clock with a master clock (an NTP server). For example, the Wireless Controller resets its clock to 07:04:59 upon reading a time of 07:04:59 from its designated NTP server.

Additionally, if the profile is supporting an Access Point, the profile's general configuration provides an option to disable the device's LEDs.

To define a profile's general configuration:

1. Select the Configuration tab from the Web UI.
2. Select **Profiles** from the Configuration tab.
3. Select **Manage Profiles** from the Configuration > Profiles menu.
4. Select **General**.

A General configuration screen displays for the new or existing controller profile.

The screenshot shows the 'General Profile' configuration screen. It is organized into several sections:

- Settings:** IP Routing is checked.
- Auto-Provisioning Policy:** Auto-Provisioning Policy is set to '<none>'. 'Learn and save network configuration' is checked.
- AP300 Auto-Provisioning Policy:** 'Adopt Unknown APs Automatically' is checked. 'Allowed List' is set to '<none>'. 'Deny List' is set to '<none>'. There are icons for adding and deleting items next to the lists.
- Network Time Protocol (NTP):** A table with columns: Server IP, Authentication Key, Prefer, Autokey, Key, and Version. An 'Add Row' button is below the table.

At the bottom of the screen are 'OK', 'Reset', and 'Exit' buttons.

FIGURE 223 General Profile - screen

5. If creating a new profile, provide a name (up to 32 characters) within the **Profile** parameter field.
6. Use the **Type** drop-down menu to specify the Brocade Mobility AP or controller model for which the profile applies.

Profiles can only be applied to the same device type selected when the profile is initially created.

7. In the **Settings** section check the **IP Routing** checkbox to enable routing for the device.
8. Refer to the **Provisioning Policy** section to select a **Provisioning Policy** or create a new one.

Provisioning Policy

Select a Provisioning Policy from the pulldown menu. To create a new Provisioning Policy click the create icon. For more information on creating a provisioning policy that can be applied to a controller profile, see *Auto Provisioning Policies* on page 5-216.

Learn and save network configuration

Check the Learn and save network configuration checkbox to enable the device to learn and save network information.

9. Select **+ Add Row** below the **Network Time Protocol (NTP)** table to define the configurations of NTP server resources the controller uses it obtain system time. Set the following parameters to define the NTP configuration:

Server IP	Set the IP address of each server added as a potential NTP resource.
Authentication Key	Select the number of the associated Authentication Key for the NTP resource.
Prefer	Select the check box to designate this particular NTP resource as preferred. If using multiple NTP resources, preferred resources will be given first opportunity to connect to the controller and provide NTP calibration.
AutoKey	Select the check box to enable an autokey configuration for the controller and NTP resource. The default setting is disabled.
Key	If an autokey is not being used, manually enter a 64 character maximum key the controller and NTP resource share to securely interoperate.
Version	Use the spinner control to specify the version number used by this NTP server resource. The default setting is 0.

10. Select **OK** to save the changes made to the general profile configuration. Select **Reset** to revert to the last saved configuration.

General Profile Configuration and Deployment Considerations

General Profile Configuration

Before defining a general profile configuration, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- A default profile is applied to the controller automatically, and default AP profiles are applied to APs discovered by the controller.
- A central difference compared to the default-radio configurations in previous releases is that default profiles are used as pointers of an AP's configuration, not just templates from which the configuration is copied. Therefore, even after adoption, if a change is made in one of the parameters in a profile, the change is reflected across all APs using that profile.
- Each user defined profile requires a unique name.
- User defined Profiles can be automatically assigned to Access Points using AP adoption policies.
- All Wireless Controllers and Access Points are automatically assigned a default profile based on the hardware type selected when the profile is initially created.

Profile Cluster Configuration (Controllers Only)

Configuration and network monitoring are two tasks a network administrator faces as a network grows in terms of the number of managed nodes (controllers, routers, wireless devices etc.). Such scalability requirements lead network administrators to look for managing and monitoring each node from a single centralized management entity. The controller not only provides a centralized management solution, it provides a centralized management profile that can be shared by any single controller in the cluster. This eliminates dedicating a management entity to manage all cluster members and eliminates a single point of failure.

A redundancy group (cluster) is a set of controllers (nodes) uniquely defined by the controller profile's configuration. Within the redundancy group, members discover and establish connections to other controller members and provide wireless network self-healing support in the event of cluster member failure.

NOTE

There is a limit of 2 controllers that can be configured in a cluster.

A cluster's AP load balance is typically distributed evenly amongst the controllers in the cluster. Define how often this profile is load balanced for AP radio distribution as often as you feel required, as radios can come and go and controller members can join and exit the cluster.

To define a controller's cluster configuration for use with a profile:

1. Select the Configuration tab from the Web UI.
2. Select **Profiles** from the Configuration tab.
3. Select **Manage Profiles** from the Configuration > Profiles menu.
4. Select **Cluster**.

A screen displays where the profile's cluster and AP load balancing configuration can be set.

Cluster Settings

Cluster Mode Active Standby

Cluster Name

Master Priority (1 to 255)

Handle STP Convergence

Force Configured State

Force Configured State Delay (3 to 1,800 minutes)

Cluster Member

Cluster VLAN (1 to 4,094)

FIGURE 224 Controller Profile - Cluster screen

5. Define the following **Cluster Settings** parameters to set this profile's cluster mode and deployment settings:

Cluster Mode	A member can be in either an <i>Active</i> or <i>Standby</i> mode. All active member controllers can adopt Access Points. Standby members only adopt Access Points when an active member has failed or sees an Access Point not adopted by a controller. The default cluster mode is Active and enabled for use with the controller profile.
Cluster Name	Define a name for the cluster name unique to its configuration or profile support requirements. The name cannot exceed 64 characters.
Master Priority	Set a priority value between 1 and 255 with the higher value being given higher priority. This configuration is the device's priority to become cluster master. In cluster environment one device from the cluster is elected as the cluster master. This configuration is the device's priority to become cluster master. The default value is 128.
Handle STP Convergence	Select the check box to enable <i>Spanning Tree Protocol (STP)</i> convergence for the controller. In general, this protocol is enabled in layer 2 networks to prevent network looping. Spanning Tree is a network layer protocol that ensures a loop-free topology in a mesh network of inter-connected layer 2 controllers. The spanning tree protocol disables redundant connections and uses the least costly path to maintain a connection between any two controllers in the network. If enabled, the network forwards data only after STP convergence. Enabling STP convergence delays the redundancy state machine execution until the STP convergence is completed (the standard protocol value for STP convergence is 50 seconds). Delaying the state machine is important to load balance APs at startup. The default setting is disabled.
Force Configured State	Select the check box to enable this controller to take over for an active controller member if it were to fail. A standby controller in the cluster takes over APs adopted by the failed active controller. If the failed active controller were to come back up, the active controller starts a timer based on the Auto Revert Delay interval. At the expiration of the Auto Revert Delay, the standby controller releases all adopted APs and goes back to a monitoring mode. The Auto Revert Delay timer is stopped and restarted if the active controller goes down and comes up during the Auto Revert Delay interval. The default value is disabled.
Force Configured State Delay	Specify a delay interval in either <i>Seconds</i> (1 - 1,800) or <i>Minutes</i> (1 - 30). This is the interval a standby controller waits before releasing adopted APs and goes back to a monitoring mode when an active controller becomes active again after a failure. The default interval is 5 seconds.

6. Within the **Cluster Member** field, select the **Cluster VLAN** checkbox to enable a spinner control to designate the controller VLAN where cluster members are reachable. Specify a VLAN in the range of 1 - 4094.
- Specify the IP Addresses of the VLAN's cluster members using the **IP Address** table.
7. Select an **Auto-Provisioning Policy** from the pulldown menu. To create a new Auto-Provisioning Policy click the create icon.

- Define the following **AP300** parameters:

Adopt List	Check this box to allow unconfigured AP300 Access Points to be adopted by the cluster.
Adopted List	Select an AP300 list from the pull-down menu to specify which AP300s are adopted by the cluster. If a suitable list is not present, click the create button to create a new list.
Deny List	Select an AP300 list from the pull-down menu to specify which AP300s are denied adoption by the cluster. If a suitable list is not present, click the create button to create a new list.

- Select **OK** to save the changes made to the profile's cluster configuration. Select **Reset** to revert to the last saved configuration.

Controller Cluster Profile Configuration and Deployment Considerations

Profile Cluster Configuration (Controllers Only)

Before defining a controller profile's cluster configuration, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- A controller cannot adopt more APs than the hardware capacity allow. This is important when the number of pooled AP and AAP licenses exceeds the aggregated AP and AAP capacity available on remaining controller members after a cluster member has failed. A cluster supported controller profile should be designed to ensure adequate AP and AAP capacity exists to address failure scenarios involving both APs and AAPs.
- When clustering is enabled for a profile and a controller failure occurs, AP and AAP licenses are persistent in the cluster even during controller reboots or power outages. If a cluster member failure were to occur, Brocade Mobility recommends clustering remain enabled on all remaining cluster members or the pooled member licenses will be lost.

Profile Adoption Configuration (APs Only)

- Select the Configuration tab from the Web UI.
- Select **Profiles** from the Configuration tab.
- Select **Manage Profiles** from the **Configuration > Profiles** menu.

The screenshot shows a configuration window for a provisioning policy. It includes the following elements:

- Controller Group:** A section with a 'Preferred Group' text input field.
- Controller VLAN:** A section with a 'VLAN' text input field, a checkbox, a dropdown menu showing '1', and a range '(1 to 4,094)'.
- Controller Hostnames:** A table with columns 'Host', 'Pool', and 'Routing Level'. Below the table is an 'Add Row' button.
- Bottom Right:** 'OK', 'Reset', and 'Exit' buttons.

FIGURE 225 Provisioning Policy - Rule Precedence screen

4. Set the following **Provisioning Policy** parameters:

Preferred Group

Define the Preferred Group used as optimal group of controllers for the Access Point's adoption. The name of the preferred group cannot exceed 64 characters.

VLAN

Select the checkbox to define a VLAN the Access Point's associating controller is reachable on. VLANs 0 and 4,095 are reserved and cannot be used by a controller VLAN.

Controller Hostnames

Enter Controller Hostnames as needed to define controller resources for Access Point adoption. Select + Add Row as needed to populate the table with IP Addresses or Hostnames of the controllers used as Access Point adoption resources into the managed network.

Profile Interface Configuration

A controller profile's interface configuration can be defined to support separate physical Ethernet configurations that are both unique and specific to RFS4000, RFS6000 and RFS7000 model controllers. Ports vary depending on controller platform, but controller models do have some of the same physical interfaces.

A controller requires its Virtual Interface be configured for layer 3 (IP) access or layer 3 service on a VLAN. A controller's Virtual Interface defines which IP address is associated with each VLAN ID the controller is connected to.

If the profile is configured to support an Access Point radio, an additional Radios option is available, unique to the Access Point's radio configuration.

A profile's Interface configuration process consists of the following:

- [Ethernet Port Configuration](#)
- [Virtual Interface Configuration](#)
- [Port Channel Configuration](#)
- [Access Point Radio Configuration](#)

Additionally, deployment considerations and guidelines for profile interface configurations are available for review prior to defining a configuration that could significantly impact the performance of the managed network. For more information, see *Profile Interface Deployment Considerations* on page 7-350.

Ethernet Port Configuration

[Profile Interface Configuration](#)

The ports available on a controller vary depending on the platform. The following ports are available on RFS4000, RFS6000 and RFS7000 model controllers:

- RFS4000 - ge1, ge2, ge3, ge4, ge5, up1
- RFS6000 - ge1, ge2, ge3, ge4, ge5, ge6, ge7, ge8, me1, up1
- RFS7000 - ge1, ge2, ge3, ge4, me1

GE ports are available on the RFS4000, RFS6000 and RFS7000 platforms. GE ports on the RFS4000 and RFS6000 are RJ-45 supporting 10/100/1000Mbps. GE ports on the RFS7000 can be RJ-45 or fiber ports supporting 10/100/1000Mbps.

ME ports are available on RFS6000 and RFS7000 platforms. ME ports are out-of-band management ports used to manage the controller via CLI or Web UI, even when the other ports on the controller are unreachable.

UP ports are available on RFS4000 and RFS6000 platforms. An UP port is used to connect the controller to the backbone network. An UP port supports either RJ-45 or fiber. The UP port is the preferred means to connect to the backbone as it has a non-blocking 1gbps connection unlike GE ports.

To define a controller profile's Ethernet port configuration:

1. Select **Configuration > Profiles > Interface**.
2. Expand the Interface menu to display its submenu options.
3. Select **Ethernet Ports**.

The Ethernet Ports screen displays configuration, runtime status and statistics regarding the physical ports on the controller.

Native VLAN

Lists the numerical VLAN ID (1 - 4094) set for the native VLAN. The native VLAN allows an Ethernet device to associate untagged frames to a VLAN when no 802.1Q frame is included in the frame. Additionally, the native VLAN is the VLAN untagged traffic is directed over when using a port in trunk mode.

Tag Native VLAN

A green checkmark defines the native VLAN as tagged. A red "X" defines the native VLAN as untagged. When a frame is tagged, the 12 bit frame VLAN ID is added to the 802.1Q header so upstream Ethernet devices know which VLAN ID the frame belongs to. The device reads the 12 bit VLAN ID and forwards the frame to the appropriate VLAN. When a frame is received with no 802.1Q header, the upstream device classifies the frame using the default or native VLAN assigned to the Trunk port. A native VLAN allows an Ethernet device to associate untagged frames to a VLAN when no 802.1Q frame is included in the frame.

Allowed VLANs

Displays those VLANs allowed to send packets over the listed controller port. Allowed VLANs are only listed when the mode has been set to Trunk.

- To edit the configuration of an existing controller port, select it from amongst those displayed and select the **Edit** button. The Ethernet port **Basic Configuration** screen displays by default.

FIGURE 227 Ethernet Ports - Basic Configuration screen

- Set the following Ethernet port **Properties**:

Description

Enter a brief description for the controller port (64 characters maximum). The description should reflect the port's intended function to differentiate it from others with similar configurations or perhaps just the name of the physical port.

Admin Status	Select the Enabled radio button to define this port as active to the controller profile it supports. Select the Disabled radio button to disable this physical controller port in the controller profile. It can be activated at any future time when needed.
Speed	Select the speed at which the port can receive and transmit the data. Select either 10 Mbps, 100 Mbps, 1000 Mbps. Select either of these options to establish a 10, 100 or 1000 Mbps data transfer rate for the selected half duplex or full duplex transmission over the port. These options are not available if Auto is selected. Select Automatic to enable the controller port to automatically exchange information about data transmission speed and duplex capabilities. Auto negotiation is helpful when in an environment where different devices are connected and disconnected on a regular basis. Automatic is the default setting.
Duplex	Select either half, full or automatic as the duplex option. Select Half duplex to send data over the port, then immediately receive data from the same direction in which the data was transmitted. Like a full-duplex transmission, a half-duplex transmission can carry data in both directions, just not at the same time. Select Full duplex to transmit data to and from the controller port at the same time. Using Full duplex, the port can send data while receiving data as well. Select Automatic to dynamically duplex as port performance needs dictate. Automatic is the default setting.

7. Enable or disable the following **CDP/LLDP** parameters used to configure Cisco Discovery Protocol and Link Layer Discovery Protocol for this profile's Ethernet port configuration:

Cisco Discovery Protocol Receive	Select this box to allow the Cisco discovery protocol to be received on this controller port. If enabled, the port sends out periodic interface updates to a multicast address to advertise its presence to neighbors. This option is enabled by default.
Cisco Discovery Protocol Transmit	Select this box to allow the Cisco discovery protocol to be transmitted on this controller port. If enabled, the port sends out periodic interface updates to a multicast address to advertise its presence to neighbors.
Link Layer Discovery Protocol Receive	Select this box to allow the Link Layer discovery protocol to be received on this controller port. If enabled, the port sends out periodic interface updates to a multicast address to advertise its presence to neighbors. This option is enabled by default.
Link Layer Discovery Protocol Transmit	Select this box to allow the Link Layer discovery protocol to be transmitted on this controller port. If enabled, the port sends out periodic interface updates to a multicast address to advertise its presence to neighbors.

Set the following **Power Over Ethernet (PoE)** parameters for this profile's Ethernet port configuration:

Enable POE	Select the check box to configure the selected port to use Power over Ethernet. To disable PoE on a port, uncheck this option. Power over Ethernet is supported on RFS4000 and RFS6000 model controllers only. When enabled, the controller supports 802.3af PoE on each of its ge ports. The PoE allows users to monitor port power consumption and configure power usage limits and priorities for each ge port.
Power Limit	Use the spinner control to set the total watts available for Power over Ethernet on the defined controller ge port. Set a value between 0 - 40 watts.
Power Priority	Set the power priority for the listed port to either to either <i>Low</i> , <i>Medium</i> or <i>High</i> . This is the priority assigned to this port versus the power requirements of the other ports on the controller.

8. Define the following **Switching Mode** parameters to apply to the Ethernet port configuration:

- Mode** Select either the *Access* or *Trunk* radio button to set the VLAN switching mode over the port. If *Access* is selected, the port accepts packets only from the native VLANs. Frames are forwarded out the port untagged with no 802.1Q header. All frames received on the port are expected as untagged and are mapped to the native VLAN. If the mode is set to *Trunk*, the port allows packets from a list of VLANs you add to the trunk. A port configured as *Trunk* supports multiple 802.1Q tagged VLANs and one Native VLAN which can be tagged or untagged. *Access* is the default mode.
- Native VLAN** Use the spinner control to define a numerical **Native VLAN ID** between 1 - 4094. The native VLAN allows an Ethernet device to associate untagged frames to a VLAN when no 802.1Q frame is included in the frame. Additionally, the native VLAN is the VLAN which untagged traffic will be directed over when using a port in trunk mode. The default VLAN is 1.
- Tag Native VLAN** Select the check box to tag the native VLAN. Controllers support the IEEE 802.1Q specification for tagging frames and coordinating VLANs between devices. IEEE 802.1Q adds four bytes to each frame identifying the VLAN ID for upstream devices that the frame belongs. If the upstream Ethernet device does not support IEEE 802.1Q tagging, it does not interpret the tagged frames. When VLAN tagging is required between devices, both devices must support tagging and be configured to accept tagged VLANs. When a frame is tagged, the 12 bit frame VLAN ID is added to the 802.1Q header so upstream Ethernet devices know which VLAN ID the frame belongs to. The device reads the 12 bit VLAN ID and forwards the frame to the appropriate VLAN. When a frame is received with no 802.1Q header, the upstream device classifies the frame using the default or native VLAN assigned to the Trunk port. The native VLAN allows an Ethernet device to associate untagged frames to a VLAN when no 802.1Q frame is included in the frame. This feature is disabled by default.
- Allowed VLANs** Selecting *Trunk* as the mode enables the **Allowed VLANs** parameter. Add VLANs that exclusively send packets over the listed port.

9. Optionally select the **Port Channel** checkbox and define a setting between 1 - 3 using the spinner control. This sets the channel group for the port.
10. Select **OK** to save the changes made to the Ethernet Port Basic Configuration. Select **Reset** to revert to the last saved configuration.
11. Select the **Security** tab.

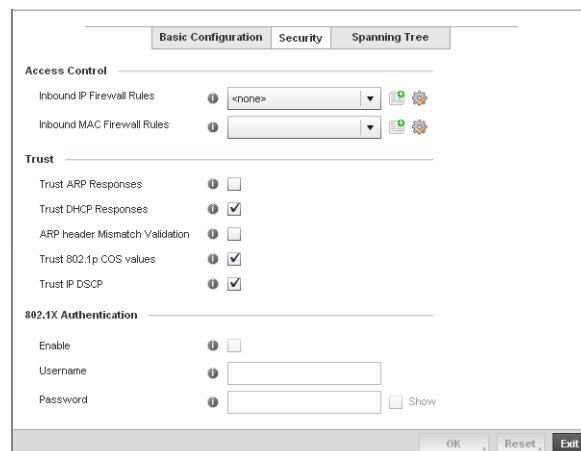


FIGURE 228 Ethernet Ports - Security screen

12. Refer to the **Access Control** field. As part of the port's security configuration, Inbound IP and MAC address firewall rules are required.

Use the **Inbound IP Firewall Rules** and **Inbound MAC Firewall Rules** drop-down menus to select the firewall rules to apply to this profile's Ethernet port configuration.

The firewall inspects IP and MAC traffic flows and detects attacks typically not visible to traditional wired firewall appliances.

13. If a firewall rule does not exist suiting the data protection needs of the target port configuration, select the **Create** icon to define a new rule configuration. For more information, see [Wireless Firewall on page 9-419](#).

14. Refer to the **Trust** field to define the following:

Trust ARP Responses	Select the check box to enable ARP trust on this controller port. ARP packets received on this controller port are considered trusted and information from these packets is used to identify rogue devices within the managed network. The default value is disabled.
Trust DHCP Responses	Select the check box to enable DHCP trust on this port. If enabled, only DHCP responses are trusted and forwarded on this port, and a DHCP server can be connected only to a DHCP trusted port. The default value is enabled.
ARP header Mismatch Validation	Select the check box to enable a mismatch check for the source MAC in both the ARP and Ethernet header. The default value is enabled.
Trust 802.1p COS values	Select the check box to enable 802.1p COS values on this port. The default value is enabled.
Trust IP DSCP	Select the check box to enable IP DSCP values on this port. The default value is enabled.

NOTE

Some vendor solutions with VRRP enabled send ARP packets with Ethernet SMAC as a physical MAC and inner ARP SMAC as VRRP MAC. If this configuration is enabled, a packet is allowed, despite a conflict existing.

15. Select **OK** to save the changes made to the Ethernet port's security configuration. Select **Reset** to revert to the last saved configuration.

16. Select the **Spanning Tree** tab.

FIGURE 229 Ethernet Ports - Spanning Tree screen

17. Define the following **PortFast** parameters for the port's MSTP configuration:

- Enable PortFast** Select the check box to enable drop-down menus for both the Enable Portfast BPDU Filter and Enable Portfast BPDU guard options for the controller port.
- PortFast BPDU Filter** Select enable to invoke a BPDU filter for this portfast enabled port. Enabling the BPDU filter feature ensures this PortFast enabled port does not transmit or receive BPDUs.
- PortFast BPDU Guard** Select enable to invoke a BPDU guard for this portfast enabled port. Enabling the BPDU Guard feature means this portfast-enabled port will shutdown on receiving a BPDU. Thus, no BPDUs are processed.

18. Set the following **MSTP Configuration** parameters:

- Enable as Edge Port** Select the check box to define this port as an edge port. Using an edge (private) port, you can isolate devices to prevent connectivity over this port.
- Link Type** Select either the *Point-to-Point* or *Shared* radio button. Selecting *Point-to-Point* indicates the port should be treated as connected to a point-to-point link. Selecting *Shared* indicates this port should be treated as having a shared connection. A port connected to a hub is on a shared link, while one connected to a controller is a point-to-point link.
- Cisco MSTP Interoperability** Select either the *Enable* or *Disable* radio buttons. This enables interoperability with Cisco's version of MSTP over the port, which is incompatible with standard MSTP.
- Force Protocol Version** Sets the protocol version to either *STP(0)*, *Not Supported(1)*, *RSTP(2)* or *MSTP(3)*. MSTP is the default setting.
- Guard** Determines whether the port enforces root bridge placement. Setting the guard to *Root* ensures the port is a designated port. Typically, each guard root port is a designated port, unless two or more ports (within the root bridge) are connected together. If the bridge receives superior (BPDUs) on a guard root-enabled port, the guard root moves the port to a root-inconsistent STP state. This state is equivalent to a listening state. No data is forwarded across the port. Thus, the guard root enforces the root bridge position.

19. Refer to the **Spanning Tree Port Cost** table.

Define an **Instance Index** using the spinner control, then set the **Cost**. The default path cost depends on the speed of the port. The cost helps determine the role of the port in the MSTP network. The designated cost is the cost for a packet to travel from this port to the root in the MSTP configuration. The slower the media, the higher the cost.

Speed	Default Path Cost
<=100000 bits/sec	200000000
<=1000000 bits/sec	20000000
<=10000000 bits/sec	2000000
<=100000000 bits/sec	200000
<=1000000000 bits/sec	20000
<=10000000000 bits/sec	2000
<=100000000000 bits/sec	200
<=1000000000000 bits/sec	20
>1000000000000 bits/sec	2

20. Select **+ Add Row** as needed to include additional indexes.

21. Refer to the **Spanning Tree Port Priority** table.

Define or override an **Instance Index** using the spinner control and then set the **Priority**. The lower the priority, a greater likelihood of the port becoming a designated port. Thus applying an higher override value impacts the port's likelihood of becoming a designated port.

Select **+ Add Row** needed to include additional indexes.

22. Select **OK** to save the changes made to the Ethernet Port Spanning Tree configuration. Select **Reset** to revert to the last saved configuration.

Virtual Interface Configuration

Profile Interface Configuration

A controller Virtual Interface is required for layer 3 (IP) access to the controller or provide layer 3 service on a VLAN. The Virtual Interface defines which IP address is associated with each VLAN ID the controller is connected to. A Virtual Interface is created for the default VLAN (VLAN 1) to enable remote controller administration. A Virtual Interface is also used to map VLANs to IP address ranges. This mapping determines the destination networks for controller routing.

To review existing Virtual Interface configurations and either create a new Virtual Interface configuration, modify an existing configuration or delete an existing configuration:

1. Select **Configuration > Profiles > Interface**.
2. Expand the Interface menu to display its submenu options.
3. Select **Virtual Interfaces**.

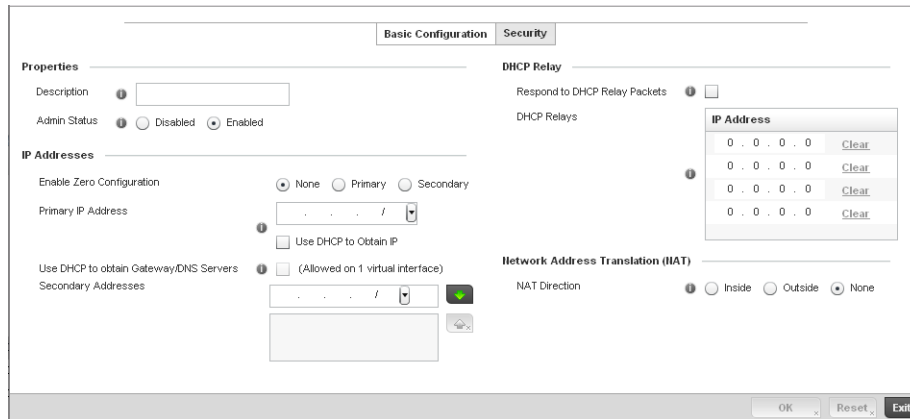


FIGURE 231 Virtual Interfaces - Basic Configuration screen

The **Basic Configuration** screen displays by default, regardless of whether a new Virtual Interface is created or an existing one is being modified.

6. If creating a new Virtual Interface, use the **VLAN ID** spinner control to define a numeric ID between 1 - 4094.
7. Define the following parameters from within the **Properties** field:

Description	Provide or edit a description (up to 64 characters) for the Virtual Interface that helps differentiate it from others with similar configurations.
Admin Status	Either select the Disabled or Enabled radio button to define this interface's current status within the managed network. When set to Enabled, the Virtual Interface is operational and available to the controller. The default value is disabled.

8. Set the following network information from within the **IP Addresses** field:

Enable Zero Configuration	Define the IP address for the VLAN associated Virtual Interface.
Primary IP Address	Define the IP address for the VLAN associated Virtual Interface.
Use DHCP to Obtain IP	Select this option to allow DHCP to provide the IP address for the Virtual Interface. Selecting this option disables the Primary IP address field.
Use DHCP to obtain Gateway/DNS Servers	Select this option to allow DHCP to obtain a default gateway address, and DNS resource for one virtual interface. This setting is disabled by default and only available when the Use DHCP to Obtain IP option is selected.
Secondary Addresses	Use the Secondary Addresses parameter to define additional IP addresses to associate with VLAN IDs. The address provided in this field is used if the primary IP address is unreachable.

9. Refer to the **DHCP Relay** field to set or override the DHCP relay server configuration used with the controller Virtual Interface.

Respond to DHCP Relay Packets Select the Respond to DHCP Relay Packets option to allow the controller's onboard DHCP server to respond to relayed DHCP packets on this interface.

DHCP Relay IP Address Provide IP addresses for DHCP server relay resources.
The interface VLAN and gateway should have their IP addresses set. The interface VLAN and gateway interface should not have DHCP client or DHCP Server enabled. DHCP packets cannot be relayed to an onboard DHCP Server. The interface VLAN and gateway interface cannot be the same.
When changing from a default DHCP address to a fixed IP address, set a static route first. This is critical when the controller is being accessed from a subnet not directly connected to the controller and the default route was set from DHCP.

10. Define the **Network Address Translation (NAT)** direction.

Select either **Inside**, **Outside** or **None**.

- *Inside* - The inside network is transmitting data over the network its intended destination. On the way out, the source IP address is changed in the header and replaced by the (public) IP address.
- *Outside* - Packets passing through the NAT on the way back to the managed LAN are searched against the records kept by the NAT engine. There, the destination IP address is changed back to the specific internal private class IP address in order to reach the LAN over the network.
- *None* - No NAT activity takes place. This is the default setting.

NOTE

Refer to *Setting the Profile's NAT Configuration on page 7-383* for instructions on creating a profile's NAT configuration.

11. Select **OK** button to save the changes to the Basic Configuration screen. Select **Reset** to revert to the last saved configuration.

12. Select the **Security** tab.

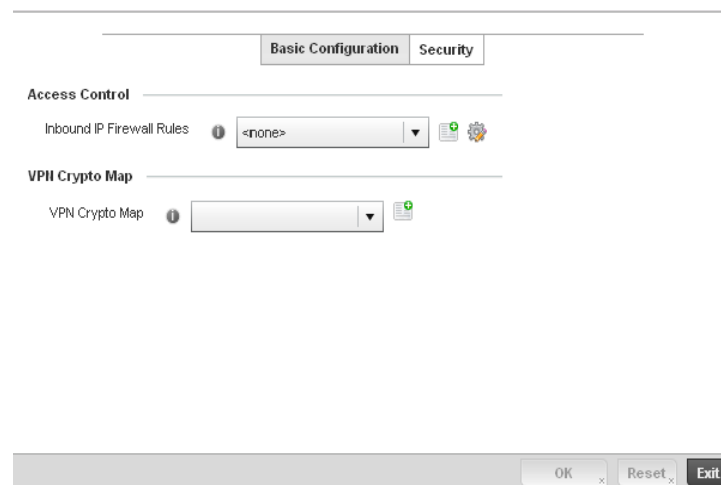


FIGURE 232 Virtual Interfaces - Security screen

13. Use the **Inbound IP Firewall Rules** drop-down menu to select the firewall rule configuration to apply to this Virtual Interface.

The firewall inspects and packet traffic to and from connected clients.

If a firewall rule does not exist suiting the data protection needs of this Virtual Interface, select the **Create** icon to define a new firewall rule configuration or the **Edit** icon to modify an existing configuration. For more information, see [Wireless Firewall on page 9-419](#).

14. Use the **VPN Crypto Map** drop-down menu to select the Crypto Map configuration to apply to this Virtual Interface.

Crypto Map entries are sets of configuration parameters for encrypting packets that pass through the VPN Tunnel. If a Crypto Map configuration does not exist suiting the needs of this Virtual Interface, select the **Create** icon to define a new Crypto Map configuration or the **Edit** icon to modify an existing configuration. For more information, see *Setting the Profile's VPN Configuration on page 7-377*.

15. Select the **OK** button located at the bottom right of the screen to save the changes to the Security screen. Select **Reset** to revert to the last saved configuration.

Port Channel Configuration

[Profile Interface Configuration](#)

Controller profiles can be applied customized port channel configurations as part of their Interface configuration.

To define a port channel configuration for a controller profile:

1. Select **Configuration > Profiles > Interface**.
2. Expand the Interface menu to display its submenu options.
3. Select **Port Channels**.

The Port Channels screen displays.

FIGURE 234 Port Channels - Basic Configuration screen

6. Set the following port channel **Properties**:

Description	Enter a brief description for the controller port channel (64 characters maximum). The description should reflect the port channel's intended function.
Admin Status	Select the Enabled radio button to define this port channel as active to the controller profile it supports. Select the Disabled radio button to disable this port channel configuration within the controller profile. It can be activated at any future time when needed. The default setting is disabled.
Speed	Select the speed at which the port channel can receive and transmit the data. Select either <i>10 Mbps</i> , <i>100 Mbps</i> , <i>1000 Mbps</i> . Select either of these options to establish a 10, 100 or 1000 Mbps data transfer rate for the selected half duplex or full duplex transmission over the port. These options are not available if <i>Auto</i> is selected. Select <i>Automatic</i> to enable the port channel to automatically exchange information about data transmission speed and duplex capabilities. Auto negotiation is helpful when in an environment where different devices are connected and disconnected on a regular basis. Automatic is the default setting.
Duplex	Select either <i>Half</i> , <i>Full</i> or <i>Automatic</i> as the duplex option. Select Half duplex to send data over the port channel, then immediately receive data from the same direction in which the data was transmitted. Like a Full duplex transmission, a Half duplex transmission can carry data in both directions, just not at the same time. Select Full duplex to transmit data to and from the port channel at the same time. Using Full duplex, the port channel can send data while receiving data as well. Select Automatic to enable to the controller to dynamically duplex as port channel performance needs dictate. Automatic is the default setting.

7. Use the **Port Channel Load Balance** drop-down menu to define whether port channel load balancing is conducted using a *Source/Destination IP* or a *Source/Destination MAC*. Source/Destination IP is the default setting.

8. Define the following **Switching Mode** parameters to apply to the port channel configuration:

Mode	Select either the <i>Access</i> or <i>Trunk</i> radio button to set the VLAN switching mode over the port channel. If <i>Access</i> is selected, the port channel accepts packets only from the native VLANs. Frames are forwarded out the port untagged with no 802.1Q header. All frames received on the port are expected as untagged and are mapped to the native VLAN. If the mode is set to <i>Trunk</i> , the port channel allows packets from a list of VLANs you add to the trunk. A port channel configured as <i>Trunk</i> supports multiple 802.1Q tagged VLANs and one Native VLAN which can be tagged or untagged. <i>Access</i> is the default setting.
Native VLAN	Use the spinner control to define a numerical ID between 1 - 4094. The native VLAN allows an Ethernet device to associate untagged frames to a VLAN when no 802.1Q frame is included in the frame. Additionally, the native VLAN is the VLAN which untagged traffic will be directed over when using trunk mode. The default value is 1.
Tag the Native VLAN	Select the checkbox to tag the native VLAN. Controllers support the IEEE 802.1Q specification for tagging frames and coordinating VLANs between devices. IEEE 802.1Q adds four bytes to each frame identifying the VLAN ID for upstream devices that the frame belongs. If the upstream Ethernet device does not support IEEE 802.1Q tagging, it does not interpret the tagged frames. When VLAN tagging is required between devices, both devices must support tagging and be configured to accept tagged VLANs. When a frame is tagged, the 12 bit frame VLAN ID is added to the 802.1Q header so upstream Ethernet devices know which VLAN ID the frame belongs to. The device reads the 12 bit VLAN ID and forwards the frame to the appropriate VLAN. When a frame is received with no 802.1Q header, the upstream device classifies the frame using the default or native VLAN assigned to the Trunk port. The native VLAN allows an Ethernet device to associate untagged frames to a VLAN when no 802.1Q frame is included in the frame. This setting is disabled by default.
Allowed VLANs	Selecting <i>Trunk</i> as the mode enables the Allowed VLANs parameter. Add VLANs that exclusively send packets over the port channel.

9. Select **OK** to save the changes made to the port channel Basic Configuration. Select **Reset** to revert to the last saved configuration.

10. Select the **Security** tab.

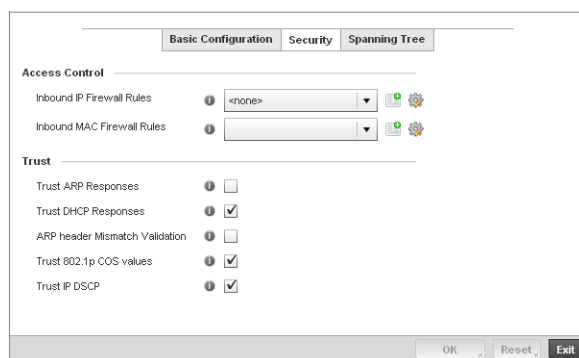


FIGURE 235 Port Channels - Security screen

11. Refer to the **Access Control** section. As part of the port channel's security configuration, Inbound IP and MAC address firewall rules are required.

Use the **Inbound IP Firewall Rules** and **Inbound MAC Firewall Rules** drop-down menus to select firewall rules to apply to this profile's port channel configuration.

The firewall inspects IP and MAC traffic flows and detects attacks typically not visible to traditional wired firewall appliances.

If a firewall rule does not exist suiting the data protection needs of the target port channel configuration, select the **Create** icon to define a new rule configuration or the **Edit** icon to modify an existing firewall rule configuration. For more information, see *Wireless Firewall on page 9-419*.

12. Refer to the **Trust** field to define the following:

- Trust ARP Responses** Select the check box to enable ARP trust on this port channel. ARP packets received on this controller port are considered trusted and information from these packets is used to identify rogue devices within the managed network. The default value is disabled.
- Trust DHCP Responses** Select the check box to enable DHCP trust. If enabled, only DHCP responses are trusted and forwarded on this port channel, and a DHCP server can be connected only to a DHCP trusted port. The default value is enabled.
- ARP header Mismatch Validation** Select the check box to enable a mismatch check for the source MAC in both the ARP and Ethernet header. The default value is enabled.
- Trust 802.1p COS values** Select the check box to enable 802.1p COS values on this port channel. The default value is enabled.
- Trust IP DSCP** Select the check box to enable IP DSCP values on this port channel. The default value is disabled.

13. Select **OK** to save the changes to the security configuration. Select **Reset** to revert to the last saved configuration.

14. Select the **Spanning Tree** tab.

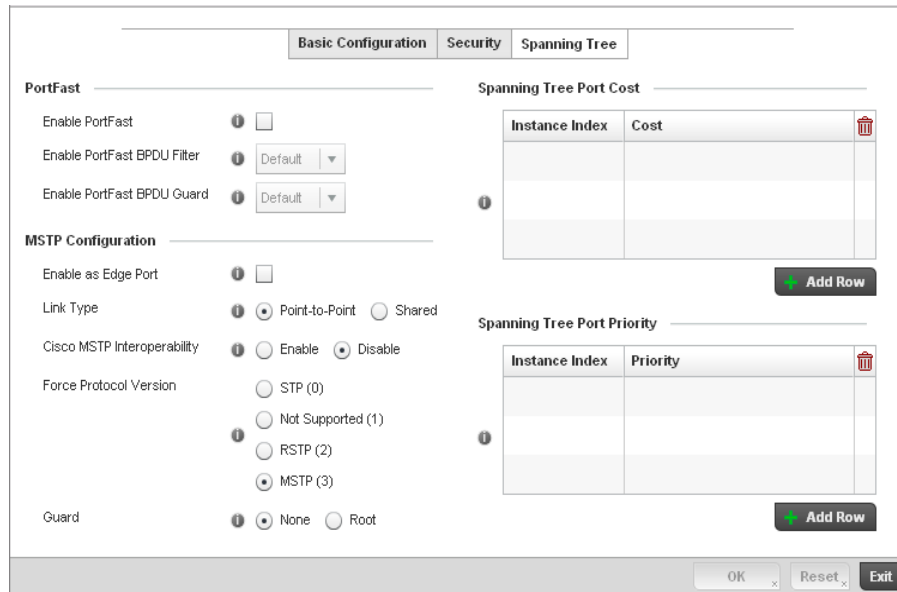


FIGURE 236 Port Channels - Spanning Tree screen

15. Define the following **PortFast** parameters for the port channel's MSTP configuration:

- Enable PortFast** Select the check box to enable drop-down menus for both the port Enable Portfast BPDU Filter and Enable Portfast BPDU guard options. This setting is disabled by default.
- PortFast BPDU Filter** Select Enable to invoke a BPDU filter for this portfast enabled port channel. Enabling the BPDU filter feature ensures this port channel does not transmit or receive any BPDUs. The default setting is None.
- PortFast BPDU Guard** Select Enable to invoke a BPDU guard for this portfast enabled port channel. Enabling the BPDU Guard feature means this port will shutdown on receiving a BPDU. Thus, no BPDUs are processed. The default setting is None.

16. Set the following **MSTP Configuration** parameters for the port channel:

- Enable as Edge Port** Select the check box to define this port as an edge port. Using an edge (private) port, you can isolate devices to prevent connectivity over this port channel. This setting is disabled by default.
- Link Type** Select either the *Point-to-Point* or *Shared* radio button. Selecting Point-to-Point indicates the port should be treated as connected to a point-to-point link. Selecting Shared indicates this port should be treated as having a shared connection. A port connected to a hub is on a shared link, while one connected to a controller is a point-to-point link. Point-to-Point is the default setting.
- Cisco MSTP Interoperability** Select either the *Enable* or *Disable* radio buttons. This enables interoperability with Cisco's version of MSTP, which is incompatible with standard MSTP. This setting is disabled by default.
- Force Protocol Version** Sets the protocol version to either *STP(0)*, *Not Supported(1)*, *RSTP(2)* or *MSTP(3)*. MSTP is the default setting.
- Guard** Determines whether the port channel enforces root bridge placement. Setting the guard to **Root** ensures the port is a designated port. Typically, each guard root port is a designated port, unless two or more ports (within the root bridge) are connected together. If the bridge receives superior (BPDUs) on a guard root-enabled port, the guard root moves the port to a root-inconsistent STP state. This state is equivalent to a listening state. No data is forwarded across the port. Thus, the guard root enforces the root bridge position.

17. Refer to the **Spanning Tree Port Cost** table.

18. Define an **Instance Index** using the spinner control and then set the cost. The default path cost depends on the user defined port speed. The cost helps determine the role of the port channel in the MSTP network. The designated cost is the cost for a packet to travel from this port to the root in the MSTP configuration. The slower the media, the higher the cost.

Speed	Default Path Cost
<=100000 bits/sec	200000000
<=1000000 bits/sec	20000000
<=10000000 bits/sec	2000000
<=100000000 bits/sec	200000
<=1000000000 bits/sec	20000

Speed	Default Path Cost
<=100000000000 bits/sec	2000
<=100000000000 bits/sec	200
<=1000000000000 bits/sec	20
>1000000000000 bits/sec	2

19. Select **+ Add Row** as needed to include additional indexes.
20. Refer to the **Spanning Tree Port Priority** table.
21. Define or override an **Instance Index** using the spinner control and then set the **Priority**. The lower the priority, a greater likelihood of the port becoming a designated port.
22. Select **+ Add Row** needed to include additional indexes.
23. Select **OK** to save the changes made to the Ethernet Port Spanning Tree configuration. Select **Reset** to revert to the last saved configuration.

Access Point Radio Configuration

Profile Interface Configuration

Access Points can have their radio configurations modified by their connected controller once their radios have successfully associated. Take care not to modify an Access Point's configuration using its resident Web UI, CLI or SNMP interfaces when managed by a controller profile, or risk the Access Point having a configuration independent from the profile until the profile can be uploaded to the Access Point again.

To define a Access Point radio configuration from the Access Point's associated controller:

1. Select **Configuration > Profiles > Interface**.
2. Expand the Interface menu to display its submenu options.
3. Select **Radios**.

5. If required, select a radio configuration and select the **Edit** button to modify its configuration.

FIGURE 238 Access Point Radio - Settings tab

6. The **Radio Settings** tab displays by default.
7. Define the following radio configuration parameters from within the **Properties** field:

Description

Provide or edit a description (1 - 64 characters in length) for the radio that helps differentiate it from others with similar configurations.

Admin Status

Either select *Active* or *Shutdown* to define this radio's current status within the managed network. When defined as Active, the Access Point is operational and available for client support within the managed network.

Radio QoS Policy

Use the drop-down menu to specify an existing QoS policy to apply to the Access Point radio in respect to its intended radio traffic. If there's no existing suiting the radio's intended operation, select the **Create** icon to define a new QoS policy that can be applied to this controller profile. For more information, see *Radio QoS Policy on page 6-278*.

Association ACL

Use the drop-down menu to specify an existing Association ACL policy to apply to the Access Point radio. An Association ACL is a policy-based ACL that either prevents or allows wireless clients from connecting to a managed Access Point radio. An ACL is a sequential collection of permit and deny conditions that apply to controller packets. When a packet is received on an interface, the controller compares the fields in the packet against any applied ACLs to verify the packet has the required permissions to be forwarded, based on the criteria specified in the access lists. If a packet does not meet any of the criteria specified in the ACL, the packet is dropped. Select the **Create** icon to define a new Association ACL that can be applied to this controller profile. For more information, see *Association ACL on page 6-298*.

8. Set the following profile **Radio Settings** for the selected Access Point radio.

RF Mode	Set the mode to either <i>2.4 GHz WLAN</i> or <i>5 GHz WLAN</i> support depending on the radio's intended client support. Set the mode to <i>Sensor</i> if using the radio for rogue device detection. To a radio as a detector, disable <i>Sensor</i> support on the other Access Point radio.
Lock RF Mode	Select the check box to lock Smart RF for this radio. The default setting is disabled.
Channel	Use the drop-down menu to select the channel of operation for the radio. Only a trained installation professional should define the radio channel. Select Smart for the radio to scan non-overlapping channels listening for beacons from other Access Points. After the channels are scanned, the radio selects the channel with the fewest Access Points. In the case of multiple Access Points on the same channel, it will select the channel with the lowest average power level. The default value is Smart.
Transmit Power	Set the transmit power of the selected Access Point radio. If using a dual or three radio model Access Point, each radio should be configured with a unique transmit power in respect to its intended client support function. A setting of 0 defines the radio as using Smart RF to determine its output power. 20 dBm is the default value.
Antenna Gain	Set the antenna between 0.00 - 15.00 dBm. The access point's <i>Power Management Antenna Configuration File</i> (PMACF) automatically configures the access point's radio transmit power based on the antenna type, its antenna gain (provided here) and the deployed country's regulatory domain restrictions. Once provided, the access point calculates the power range. Antenna gain relates the intensity of an antenna in a given direction to the intensity that would be produced ideally by an antenna that radiates equally in all directions (isotropically), and has no losses. Although the gain of an antenna is directly related to its directivity, its gain is a measure that takes into account the efficiency of the antenna as well as its directional capabilities. Brocade Mobility recommends that only a professional installer set the antenna gain. The default value is 0.00.
Antenna Mode	Set the number of transmit and receive antennas on the Access Point. 1x1 is used for transmissions over just the single "A" antenna, 1x3 is used for transmissions over the "A" antenna and all three antennas for receiving. 2x2 is used for transmissions and receipts over two antennas for dual antenna models. The default setting is dynamic based on the Access Point model deployed and its transmit power settings.
Enable Antenna Diversity	Select this box to enable antenna diversity on supported antennas. Antenna diversity uses two or more antennas to increase signal quality and strength. This option is disabled by default.
Data Rates	Once the radio band is provided, the Data Rates drop-down menu populates with rate options depending on the 2.4 or 5 GHz band selected. If the radio band is set to <i>Sensor</i> or <i>Detector</i> , the Data Rates drop-down menu is not enabled, as the rates are fixed and not user configurable. If 2.4 GHz is selected as the radio band, select separate 802.11b, 802.11g and 802.11n rates and define how they are used in combination. If 5 GHz is selected as the radio band, select separate 802.11a and 802.11n rates then define how they are used together. When using 802.11n (in either the 2.4 or 5 GHz band), Set a MCS (modulation and coding scheme) in respect to the radio's channel width and guard interval. A MCS defines (based on RF channel conditions) an optimal combination of 8 data rates, bonded channels, multiple spatial streams, different guard intervals and modulation types. Clients can associate as long as they support basic MCS (as well as non-11n basic rates).
Radio Placement	Use the drop-down menu to specify whether the radio is located <i>Indoors</i> or <i>Outdoors</i> . The placement should depend on the country of operation and its regulatory domain requirements for radio emissions. The default setting is <i>Indoors</i> .
Max Clients	Use the spinner control to set a maximum permissible number of clients to connect with this radio. The available range is between 0 - 256 clients. The default value is 256.

9. Set the following profile **WLAN Properties** for the selected Access Point radio.

Beacon Interval	Set the interval between radio beacons in milliseconds (either 50, 100 or 200). A beacon is a packet broadcast by adopted radios to keep the network synchronized. Included in a beacon is information such as the WLAN service area, the radio address, the broadcast destination addresses, a time stamp, and indicators about traffic and delivery such as a DTIM. Increase the DTIM/beacon settings (lengthening the time) to let nodes sleep longer and preserve battery life. Decrease these settings (shortening the time) to support streaming-multicast audio and video applications that are jitter-sensitive. The default value is 100 milliseconds.
DTIM Interval BSSID	Set a DTIM Interval to specify a period for <i>Delivery Traffic Indication Messages</i> (DTIM). A DTIM is periodically included in a beacon frame transmitted from adopted radios. The DTIM period determines how often the beacon contains a DTIM, for example, 1 DTIM for every 10 beacons. The DTIM indicates broadcast and multicast frames (buffered at the Access Point) are soon to arrive. These are simple data frames that require no acknowledgement, so nodes sometimes miss them. Increase the DTIM/ beacon settings (lengthening the time) to let nodes sleep longer and preserve their battery life. Decrease these settings (shortening the time) to support streaming multicast audio and video applications that are jitter-sensitive.
RTS Threshold	Specify a <i>Request To Send</i> (RTS) threshold (between 1 - 2,347 bytes) for use by the WLAN's adopted Access Point radios. RTS is a transmitting station's signal that requests a <i>Clear To Send</i> (CTS) response from a receiving client. This RTS/CTS procedure clears the air where clients are contending for transmission time. Benefits include fewer data collisions and better communication with nodes that are hard to find (or hidden) because of other active nodes in the transmission path. Control RTS/CTS by setting an RTS threshold. This setting initiates an RTS/CTS exchange for data frames larger than the threshold, and sends (without RTS/CTS) any data frames smaller than the threshold. Consider the trade-offs when setting an appropriate RTS threshold for the WLAN's Access Point radios. A lower RTS threshold causes more frequent RTS/CTS exchanges. This consumes more bandwidth because of additional latency (RTS/CTS exchanges) before transmissions can commence. A disadvantage is the reduction in data-frame throughput. An advantage is quicker system recovery from electromagnetic interference and data collisions. Environments with more wireless traffic and contention for transmission make the best use of a lower RTS threshold. A higher RTS threshold minimizes RTS/CTS exchanges, consuming less bandwidth for data transmissions. A disadvantage is less help to nodes that encounter interference and collisions. An advantage is faster data-frame throughput. Environments with less wireless traffic and contention for transmission make the best use of a higher RTS threshold.
Short Preamble	If using an 802.11bg radio, select this checkbox for the radio to transmit using a short preamble. Short preambles improve throughput. However, some devices (SpectraLink phones) require long preambles. The default value is disabled.
Guard Interval	Use the drop-down menu to specify a <i>Long</i> or <i>Any</i> guard interval. The guard interval is the space between the packets being transmitted. The guard interval is there to eliminate <i>inter-symbol interference</i> (ISI). ISI occurs when echoes or reflections from one transmission interfere with another. Adding time between transmissions allows echo's and reflections to settle before the next packet is transmitted. A shorter guard interval results in a shorter times which reduces overhead and increases data rates by up to 10%. The default value is Long.
Probe Response Rate	Use the drop-down menu to specify the data transmission rate used for the transmission of probe responses. Options include, <i>highest-basic</i> , <i>lowest-basic</i> and <i>follow-probe-request</i> (default setting).
Probe Response Retry	Select the check box to retry probe responses if they are not acknowledged by the target wireless client. The default value is enabled.

10. Select the **Enable Off Channel Scan** check box to enable scanning across all channels using this radio. Channel scans use Access Point resources and can be time consuming, so only enable when your sure the radio can afford the bandwidth be directed towards to the channel scan and does not negatively impact client support.
11. Select a mode from the **Feed WLAN Packets to Sensor** check box in the **Radio Share** section to enable this feature. Select either **Inline** or **Promiscuous** mode to allow the packets the radio is switching to also be used by the WIPS analysis module. This feature can be enabled in two modes: an inline mode where the wips sensor receives the packets from the radios with radio operating in normal mode. A promiscuous mode where the radio is configured to a mode where it receives all packets on the channel whether the destination adres is the radio or not, and the wips module can analyze them.
12. Select the **WLAN Mapping** tab.

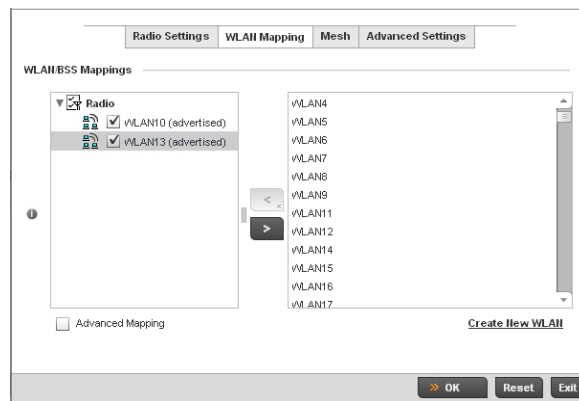


FIGURE 239 Access Point Radio - WLAN Mapping screen

13. Refer to the **WLAN/BSS Mappings** field to set WLAN BSSID assignments for an existing Access Point deployment.

Administrators can assign each WLAN its own BSSID. If using a single-radio access point, there are 8 BSSIDs available. If using a dual-radio access point there are 8 BSSIDs for the 802.11b/g/n radio and 8 BSSIDs for the 802.11a/n radio.
14. Select **Advanced Mapping** to enable WLAN mapping to a specific BSS ID.
15. Select the **OK** button located at the bottom right of the screen to save the changes to the WLAN Mapping. Select **Reset** to revert to the last saved configuration.
16. Select the **MeshConnex** tab.

FIGURE 240 Profile Overrides - Access Point Radio Mesh tab

17. Refer to the **Advanced Settings** field to define or override basic mesh settings for the Access Point radio.

Mesh

Use the pulldown to set the mesh mode for this radio. Available options are Disabled, Portal or Client. Setting the mesh mode to Disabled deactivates all mesh activity on this radio. Setting the mesh mode to Portal turns the radio into a mesh portal. This will start the radio beaconing immediately and will accept connections from other mesh nodes. Setting the mesh mode to client enables the radio to operate as a mesh client that will scan for and connect to mesh portals or nodes that are connected to portals.

Mesh Links

Specify the number of mesh links allowed by the radio. The radio can have between 1-6 mesh links when the radio is configured as a Portal or Client.

NOTE

Only single hop mesh links are supported at this time.

NOTE

The mesh encryption key is configurable from the *Command Line Interface (CLI)* using the command 'mesh_psk'. Administrators must ensure that this key is configured on the AP when it is being staged for mesh, and also added to the mesh client as well as to the portal APs configuration on the controller. For more information about the CLI please see the *Brocade Mobility RFS4000, RFS6000, and RFS7000 CLI Reference Guide*.

18. Refer to the **Preferred Peer Device** table to add mesh peers. For each peer being added enter its MAC Address and a Priority between 1 and 6. The lower the priority number the higher priority it'll be given when connecting to mesh infrastructure.
19. Select the **+ Add Row** button to add preferred peer devices for the radio to connect to in mesh mode.
20. Select the **Advanced Settings** tab.

FIGURE 241 Access Point Radio - Advanced Settings screen

21. Refer to the **Aggregate MAC Protocol Data Unit (A-MPDU)** field to define how MAC service frames are aggregated by the Access Point radio.

A-MPDU Modes

Use the drop-down menu to define the A-MPDU mode supported. Options include *Transmit Only*, *Receive Only*, *Transmit and Receive* and *None*. The default value is *Transmit and Receive*. Using the default value, long frames can be both sent and received (up to 64 KB). When enabled, define either a transmit or receive limit (or both).

Minimum Gap Between Frames

Use the drop-down menu to define the minimum gap between A-MPDU frames (in microseconds). The default value is 4 microseconds.

Received Frame Size Limit

If a support mode is enable allowing A-MPDU frames to be received, define an advertised maximum limit for received A-MPDU aggregated frames. Options include 8191, 16383, 32767 or 65535 bytes. The default value is 65535 bytes.

Transmit Frame Size Limit

Use the spinner control to set limit on transmitted A-MPDU aggregated frames. The available range is between 0 - 65,535 bytes). The default value is 65535 bytes.

22. Use the **A-MSDU Modes** drop-down menu in the **Aggregate MAC Service Data Unit (A-MSDU)** section to set or override the supported A-MSDU mode.

Available modes include *Receive Only* and *Transmit and Receive*. *Transmit and Receive* is the default value. Using *Transmit and Receive*, frames up to 4 KB can be sent and received. The buffer limit is not configurable.

23. Define a **RIFS Mode** using the drop-down menu in the **Reduced Interframe Spacing (RIFS)** section. This value determines whether interframe spacing is applied to Access Point transmissions or received packets, or both or none. The default mode is *Transmit and Receive*.

Consider setting this value to *None* for high priority traffic to reduce packet delay.

24. Set the following **Non-Unicast Traffic** values for the profile's supported Access Point radio and its connected wireless clients:

Broadcast/Multicast Transmit Rate

Use the drop-down menu to define the data rate broadcast and multicast frames are transmitted. Seven different rates are available if the not using the same rate for each BSSID, each with a separate menu.

Broadcast/Multicast Forwarding

Define whether client broadcast and multicast packets should always follow DTIM, or only follow DTIM when using Power Save Aware mode. The default setting is *Follow DTIM*.

25. Refer to the **Sniffer Redirect (Packet Capture)** field to define the radio's captured packet configuration.

Host for Redirected Packets	If packets are re-directed from a controller's connected Access Point radio, define an IP address of a resource (additional host system) used to capture the re-directed packets. This address is the numerical (non DNS) address of the host used to capture the re-directed packets.
Channel to Capture Packets	Use the drop-down menu to specify the channel used to capture re-directed packets. The default value is channel 1.

26. Select the **OK** button located at the bottom right of the screen to save the changes to the Advanced Settings screen. Select **Reset** to revert to the last saved configuration.

WAN Backhaul Override Configuration

Profile Interface Configuration

A *Wireless Wide Area Network (WWAN)* card is a specialized network interface card that allows a network device to connect, transmit and receive data over a Cellular Wide Area Network. The RFS4000, RFS6000, and certain models of the 7131 Access Point have a PCI Express card slot that supports 3G WWAN cards. The WWAN card uses point to point protocol (PPP) to connect to the Internet Service Provider (ISP) and gain access to the Internet. PPP is the protocol used for establishing internet links over dial-up modems, DSL connections, and many other types of point-to-point communications. PPP packages your system's TCP/IP packets and forwards them to the serial device where they can be put on the network. PPP is a full-duplex protocol that can be used on various physical media, including twisted pair or fiber optic lines or satellite transmission. It uses a variation of High Speed Data Link Control (HDLC) for packet encapsulation.

To define a WAN Backhaul configuration override:

1. Select **Devices** from the Configuration tab.
The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.
2. Select a target Access Point (by double-clicking it) from amongst those displayed within the Device Configuration screen.
Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.
3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Interface** to expand its sub menu options.
5. Select **WAN Backhaul**.

WAN (3G) Backhaul

WAN Interface Name *

Enable WAN (3G) Disabled Enabled

Basic Settings

Username

Password Show

Access Point Name (APN)

Authentication Type

Security Settings

VPN Crypto Map

NAT Settings

WAN NAT Direction

OK Reset Exit

FIGURE 242 Profile Overrides -WAN Backhaul screen

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

6. Refer to the **WAN (3G) Backhaul** configuration to specify WAN card settings:

WAN Interface Name	Displays the WAN Interface name for the WAN 3G Backhaul card.
Reset WAN Card	If the WAN Card becomes unresponsive or is experiencing other errors click the Reset WAN Card button to power cycle and reboot the WAN card.
Enable WAN (3G)	Check this box to enable 3G WAN card support on the device. A supported 3G card must be connected to the device for this feature to work.

7. Define or override the following authentication parameters from within the **Basic Settings** field:

Username	Provide your username for authentication support by your cellular data carrier.
Password	Provide your password for authentication support by your cellular data carrier.
Access Point Name (APN)	Enter the name of the cellular data provider if necessary. This setting is needed in areas with multiple cellular data providers using the same protocols such as Europe, the middle east and Asia.
Authentication Type	Use the pull-down menu to specify authentication type used by your cellular data provider. Supported authentication types are None, PAP, CHAP, MSCHAP, and MSCHAP-v2.

8. Select **OK** to save or override the changes to the Advanced Settings screen. Select **Reset** to revert to the last saved configuration.

Profile Interface Deployment Considerations

Profile Interface Configuration

Before defining a profile's interface configuration (supporting controller's Ethernet port, Virtual Interface, port channel and Access Point radio configurations) refer to the following deployment guidelines to ensure these configuration are optimally effective:

- Power over Ethernet is supported on RFS4000 and RFS6000 model controllers only. When enabled, the controller supports 802.3af PoE on each of its ge ports.
- When changing from a default DHCP address to a fixed IP address, set a static route first. This is critical when the controller is being accessed from a subnet not directly connected to the controller and the default route was set from DHCP.
- Take care not to modify an Access Point's configuration using its resident Web UI, CLI or SNMP interfaces when managed by a controller profile, or risk the Access Point having a configuration independent from the profile until the profile can be uploaded to the Access Point once again.

Profile Network Configuration

Setting a profile's network configuration is a large task comprised of numerous controller administration activities.

A profile's network configuration process consists of the following:

- [Setting a Profile's DNS Configuration](#)
- [ARP](#)
- [Quality of Service \(QoS\)](#)
- [Spanning Tree](#)
- [Static Routes](#)
- [Forwarding Database](#)
- [Bridge VLAN](#)
- [Miscellaneous Network Configuration](#)

Before beginning any of the profile network configuration activities described in the sections above, review the configuration and deployment considerations available in *Profile Network Configuration and Deployment Considerations* on page 7-368.

Setting a Profile's DNS Configuration

Profile Network Configuration

Domain Naming System (DNS) DNS is a hierarchical naming system for resources connected to the Internet or a private network. Primarily, the controller's DNS resources translate domain names into IP addresses. If one DNS server doesn't know how to translate a particular domain name, it asks another one until the correct IP address is returned. DNS enables access to resources using human friendly notations. DNS converts human friendly domain names into notations used by different networking equipment for locating resources.

As a resource is accessed (using human-friendly hostnames), it's possible to access the resource even if the underlying machine friendly notation name changes. Without DNS, in the simplest terms, you would need to remember a series of numbers (123.123.123.123) instead of an easy to remember domain name (for example, *www.domainname.com*).

The controller maintains its own DNS facility that can assist in domain name translation.

To define the controller's DNS configuration:

1. Select **Configuration > Profiles > Network**.
2. Expand the Network menu to display its submenu options.
3. Select **DNS**.

FIGURE 243 DNS screen

4. Set or override the following controller **Domain Name System (DNS)** configuration data:

Domain Name	Provide or override the default Domain Name used to resolve DNS names. The name cannot exceed 64 characters.
Enable Domain Lookup	Select the check box to enable DNS on the controller. When enabled, the controller can convert human friendly domain names into numerical IP destination addresses. The radio button is selected by default.
DNS Server Forwarding	Click to enable the forwarding DNS queries to external DNS servers if a DNS query cannot be processed by the controller's own DNS resources. This feature is disabled by default.

5. Set or override the following controller **DNS Server** configuration data:

Name Servers Provide a list of up to three DNS servers to forward DNS queries if the controller's DNS resources are unavailable. The DNS name servers are used to resolve IP addresses. Use the **Clear** link next to each DNS server to clear the DNS name server's IP address from the list.

6. Select **OK** to save the changes made to the DNS configuration. Select **Reset** to revert to the last saved configuration.

ARP

Profile Network Configuration

Address Resolution Protocol (ARP) is a protocol for mapping an IP address to a hardware MAC address recognized on the managed network. ARP provides protocol rules for making this correlation and providing address conversion in both directions.

When an incoming packet destined for a host arrives at the controller, the controller gateway uses ARP to find a physical host or MAC address that matches the IP address. ARP looks in its ARP cache and, if it finds the address, provides it so the packet can be converted to the right packet length and format and sent to the destination. If no entry is found for the IP address, ARP broadcasts a request packet in a special format to all the machines on the LAN to see if one machine knows that it has that IP address associated with it. A machine that recognizes the IP address as its own returns a reply so indicating. ARP updates the ARP cache for future reference and then sends the packet to the MAC address that replied.

To define an ARP supported configuration on the controller:

1. Select **Configuration > Profiles > Network**.
2. Expand the Network menu to display its submenu options.
3. Select **ARP**.
4. Select **+ Add Row** from the lower right-hand side of the screen to populate the ARP table with rows used to define ARP network address information.

Address Resolution Protocol (ARP)

Switch VLAN Interface	IP Address	MAC Address	Device Type	

Add Row

OK Reset Exit

FIGURE 244 ARP screen

- Set the following parameters to define the controller's ARP configuration:

Switch VLAN Interface	Use the spinner control to select a switch VLAN interface for an address requiring resolution.
IP Address	Define the IP address used to fetch a MAC Address.
MAC Address	Displays the target MAC address that's subject to resolution. This is the MAC used for mapping an IP address to a MAC address that's recognized on the managed network.
Device Type	Specify the device type the ARP entry supports. Host is the default setting.

- To add additional ARP overrides click on the **+ Add Row** button and enter the configuration information in the table above.
- Select the **OK** button located at the bottom right of the screen to save the changes to the ARP configuration. Select **Reset** to revert to the last saved configuration.

Quality of Service (QoS)

Profile Network Configuration

The controller uses different *Quality of Service (QoS)* screens to define WLAN and device radio QoS configurations. The controller's **Configuration > Profiles > Network** facility is separate from WLAN and radio QoS configurations, and is used to configure the priority of the different DSCP packet types.

QoS values are required to provide priority of service to some packets over others. For example, VoIP packets get higher priority than data packets to provide a better quality of service for high priority voice traffic.

The profile QoS screen maps the 6-bit *Differentiated Service Code Point* (DSCP) code points to the older

3-bit IP Precedent field located in the Type of Service byte of an IP header. DSCP is a protocol for specifying and controlling network traffic by class so that certain traffic types get precedence. DSCP specifies a specific per-hop behavior that is applied to a packet.

To define an QoS configuration for controller DSCP mappings:

1. Select **Configuration > Profiles > Network**.
2. Expand the Network menu to display its submenu options
3. Select **Quality of Service (QoS)**.

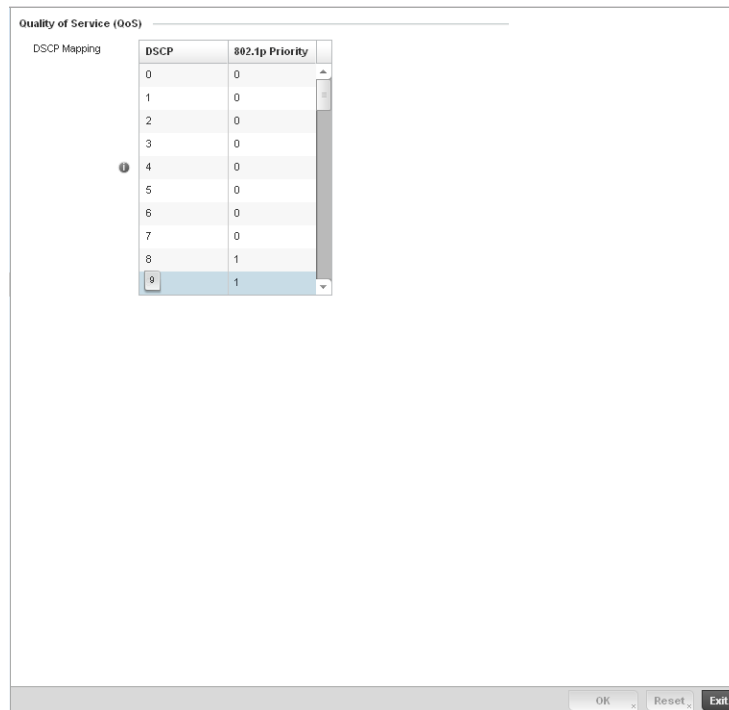


FIGURE 245 Profile QoS screen

4. Set the following parameters for IP DSCP mappings for untagged frames:

DSCP	Lists the DSCP value as a 6-bit parameter in the header of every IP packet used for packet classification.
802.1p Priority	Assign a 802.1p priority as a 3-bit IP precedence value in the Type of Service field of the IP header used to set the priority. The valid values for this field are 0-7. Up to 64 entries are permitted. The priority values are: 0 - <i>Best Effort</i> 1 - <i>Background</i> 2 - <i>Spare</i> 3 - <i>Excellent Effort</i> 4 - <i>Controlled Load</i> 5 - <i>Video</i> 6 - <i>Voice</i> 7 - <i>Network Control</i>

5. Use the spinner controls within the **802.1p Priority** field for each **DSCP** row to change its priority value.
6. Select the **OK** button located at the bottom right of the screen to save the changes. Select **Reset** to revert to the last saved configuration.

Spanning Tree

[Profile Network Configuration](#)

Spanning Tree is a network layer protocol that ensures a loop-free topology in a mesh network of inter-connected layer 2 controllers. The spanning tree protocol disables redundant connections and uses the least costly path to maintain a connection between any two controllers in the network. Spanning tree protocol allows a network design that has one or more redundant links that provide a backup path if an active link fails. This switchover is automatic and does not require any human intervention.

Physical layer redundancy may also be provided using spanning tree. Spanning tree is a link management protocol that is part of the IEEE 802.1 standard for media access control bridges. Using the Dijkstra algorithm, STP provides link path redundancy between Ethernet devices while preventing undesirable loops in a network that can be created when multiple active paths exist between Ethernet controllers and bridges.

To establish path redundancy, STP creates a tree that spans all of the controllers in an extended network, forcing redundant paths into a blocked, state. STP allows only one active path at a time between any two network devices but establishes the redundant links as a backup if the preferred link should fail. If STP costs change, or if one network segment in the STP becomes unreachable, the spanning tree algorithm reconfigures the spanning tree topology and re-establishes the link by activating the standby path. Without spanning tree, multiple paths in the Ethernet network would be active resulting in an endless loop of traffic on the LAN.

Spanning Tree is supported on the RFS4000, RFS6000 and RFS7000 model controllers and can be used to provide link path redundancy when the controllers are connected to one or more external Ethernet switches. Spanning Tree can only support one active path per VLAN between Ethernet devices. If multiple paths per VLAN exist, redundant paths are blocked.

Multiple Spanning Tree Protocol (MSTP) is a VLAN-aware protocol and algorithm to create and maintain a loop-free network. It allows the configuration of multiple spanning tree instances. This ensures a loop-free topology for one or more VLANs. It allows the administrator to define a different path for each group of VLANs to better utilize redundancy.

Using MSTP, the network can be divided into regions. Each controller within a region uses the same VLAN to instance mapping. The entire network runs a spanning tree instance called the *Common Spanning Tree* instance (CST) that interconnects regions as well as legacy (STP and RSTP) bridges. The regions run on a local instance for each configured MSTP instance.

To define a spanning tree supported configuration on the controller:

1. Select **Configuration > Profiles > Network**.
2. Expand the Network menu to display its submenu options.
3. Select **Spanning Tree**.

FIGURE 246 Spanning Tree screen

4. Set the following **MSTP Configuration** parameters:

MSTP Enable

Enables the *Multiple Spanning Tree Protocol (MSTP)* feature. Select the check box to enable spanning tree for this device. This feature is disabled by default.

Max. Hop Count

Set the maximum number of hops used when creating a Spanning Tree. This value represents the maximum allowed hops for a BPDUs (*Bridge Protocol Data Unit*) in an MSTP region. This value is used by all the MSTP instances. Enter a value between 7 - 127, or use the spinner control to set the value. The default setting is 20.

MST Config Name

Enter a name for the MST region. This is used when configuring multiple regions within the network. Each controller running MSTP is configured with a unique MST region name. This helps when keeping track of MSTP configuration changes. The name cannot exceed 64 characters.

MST Revision Level	Assign a MST revision level (0 - 255) to the MSTP region to which the device belongs. Each controller is configured with a unique MSTP name and revision number. This helps when keeping track of MSTP configuration changes. Increment this number with each configuration change. The revision level specifies the revision level of the current configuration. The default setting is 0.
Cisco MSTP Interoperability	Select <i>Enable</i> or <i>Disable</i> from the drop-down menu. This enables interoperability with Cisco's version of MSTP, which is incompatible with standard MSTP. The default setting is disabled
Hello Time	The hello time is the time interval (in seconds) the device waits between BPDU transmissions. A low value leads to excessive traffic on the network, whereas a higher value delays the detection of a topology change. Set a hello time between 1 - 10 seconds. You can also use the spinner control next to the text-box to increase or decrease the value. The default setting is 2.
Forward Delay	The forward delay is the maximum time (in seconds) the root device waits before changing states (from a listening state to a learning state to a forwarding state). Set a value between 4 -30. You can also use the spinner control next to the text-box to increase or decrease the value. The default is 15.
Maximum Age	The max-age is the maximum time (in seconds) for which, if a bridge is the root bridge, a message is considered valid. This prevents frames from looping indefinitely. The max-age should be greater than twice the value of hello time plus one, but less than twice the value of forward delay minus one. Configure this value sufficiently high, so a frame generated by root can be propagated to the leaf nodes without exceeding the max age. Set the value from 6 - 40. Use the spinner control next to the text-box to increase or decrease the value. The default setting is 20.

5. Define the following **PortFast** configuration parameters:

PortFast BPDU Filter	Select the check box to enable BPDU filter for all portfast enabled ports. The Spanning Tree Protocol sends BPDUs from all the ports. Enabling the BPDU filter feature ensures PortFast enabled ports do not transmit or receive any BPDUs.
PortFast BPDU Guard	Select the check box to enable BPDU guard for all portfast enabled ports. When the BPDU Guard feature is set for bridge, all portfast-enabled ports of the bridge that have BPDU set to default shutdown the port on receiving a BPDU. Hence no BPDUs are processed.

6. Set the following **Error Disable** recovery parameters:

Enable Recovery	Select this check box to enable an error disable timeout caused by a BPDU guard. This option is disabled by default.
Recovery Interval	Define an interval (between 10 - 1,000,000) after which a recovering port is enabled. The default recovery interval is 300.

7. Set the **Spanning Tree Instance** configuration.

Define a numerical index for each instance to assign each a unique priority. The **Priority** is assigned to an individual bridge based on whether it is selected as the root bridge. The lower the priority, the greater likelihood the bridge becoming the root for this instance.

- Use the **+ Add Row** button to create a new row in the table. To delete a row, select the row's delete icon.
- Refer to the **VLANs** table to associate a VLAN ID with the Instance index. You can add multiple VLANs to an instance.

10. Use the **+ Add Row** button to create a new row in the table. To delete a row, select the row's delete icon.
11. Select **OK** to save the changes. Select **Reset** to revert to the last saved configuration.

Static Routes

Profile Network Configuration

Use the **Static Routes** screen to set Destination IP and Gateway addresses enabling assignment of static IP addresses for requesting clients without creating numerous host pools with manual bindings. This eliminates the need for a long configuration file and reduces the controller resource space required to maintain address pools.

To create static routes:

1. Select **Configuration > Profiles > Network**.
2. Expand the Network menu to display its submenu options
3. Select **Static Routes**.

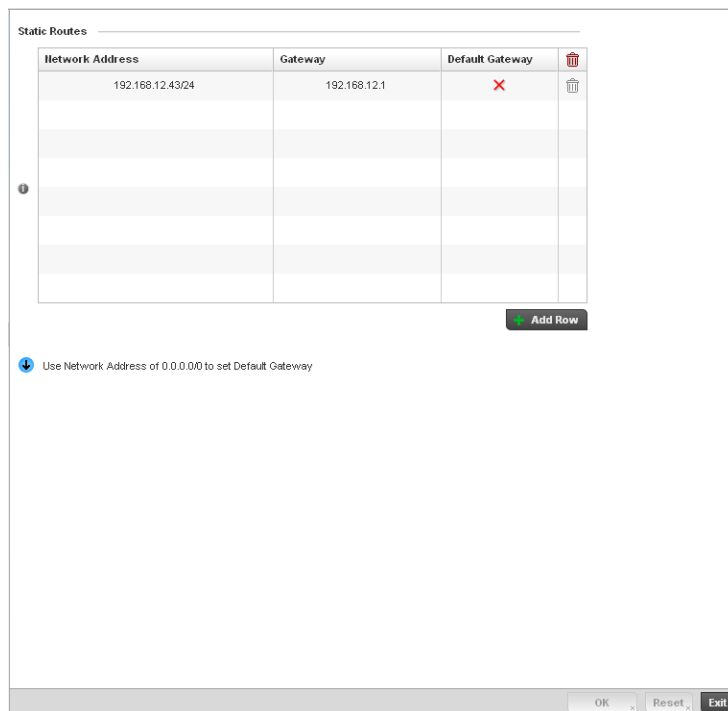


FIGURE 247 Static Routes screen

4. Select **Add Row +** as needed to include single rows in the static routes table.
5. Add **Network Addresses** and network masks in the **Network** column.
6. Provide the **Gateway** used to route traffic.
7. Select the **OK** button located at the bottom right of the screen to save the changes. Select **Reset** to revert to the last saved configuration.

Forwarding Database

Profile Network Configuration

A *Forwarding Database* is used by a bridge to forward or filter packets on behalf of the controller. The bridge reads the packet's destination MAC address and decides to either forward the packet or drop (filter) it. If it is determined the destination MAC is on a different network segment, it forwards the packet to the segment. If the destination MAC is on the same network segment, the packet is dropped (filtered). As nodes transmit packets through the bridge, the bridge updates its forwarding database with known MAC addresses and their locations on the network. This information is then used to filter or forward the packet.

To define a forwarding database configuration:

1. Select **Configuration > Profiles > Network**.
2. Expand the Network menu to display its submenu options.
3. Select **Forwarding Database**.

The screenshot shows the configuration interface for the Forwarding Database. At the top, under 'Aging Time', there is a 'Bridge Aging Time' field set to 300, with a range of (0,10-1000000 seconds). Below this is the 'Static Forwarding Table' which is a table with three columns: 'MAC Address', 'VLAN Id', and 'Interface Name'. The first row contains the values '01-02-05-08-90-48', '1', and 'EN01'. There are several empty rows below it. An 'Add Row' button is located at the bottom right of the table area. At the very bottom of the screen, there are 'OK', 'Reset', and 'Exit' buttons.

MAC Address	VLAN Id	Interface Name
01-02-05-08-90-48	1	EN01

FIGURE 248 Forwarding Database screen

4. Define a **Bridge Aging Time** between 0, 10-1,000,000 seconds.

The aging time defines the length of time an entry will remain in the a bridge's forwarding table before being deleted due to lack of activity. If an entry replenishes a destination generating continuous traffic, this timeout value will never be invoked. However, if the destination becomes idle, the timeout value represents the length of time that must be exceeded before an entry is deleted from the forwarding table. The default setting is 300 seconds.

5. Use the **+ Add Row** button to create a new row within the MAC address table.

6. Set a destination **MAC Address** address. The bridge reads the controller packet's destination MAC address and decides to forward the packet or drop (filter) it. If it's determined the destination MAC is on a different network, it forwards the packet to the segment. If the destination MAC is on the same network segment, the packet is dropped (filtered).
7. Define the target **VLAN ID** if the destination MAC is on a different network segment.
8. Provide an **Interface Name** used as the target destination interface for the target MAC address.
9. Select **OK** to save the changes. Select **Reset** to revert to the last saved configuration.

Bridge VLAN

Profile Network Configuration

A *Virtual LAN* (VLAN) is separately administrated virtual network within the same physical managed network. VLANs are broadcast domains defined to allow control of broadcast, multicast, unicast, and unknown unicast within a Layer 2 device.

Administrators often need to route traffic to interoperate between different VLANs. Bridging VLANs are only for non-routable traffic, like tagged VLAN frames destined to some other device which will untag it. When a data frame is received on a port, the VLAN bridge determines the associated VLAN based on the port of reception. Using forwarding database information, the Bridge VLAN forwards the data frame on the appropriate port(s). VLAN's are useful to set separate networks to isolate some computers from others, without actually having to have separate cabling and Ethernet switches. Controllers can do this on their own, without need for the computer or other gear to know itself what VLAN it's on (this is called port-based VLAN, since it's assigned by port of the switch). Another common use is to put specialized devices like VoIP Phones on a separate network for easier configuration, administration, security or service quality.

To define a bridge VLAN configuration:

1. Select **Configuration > Profiles > Network**.
2. Expand the Network menu to display its submenu options.
3. Select **Bridge VLAN**.

VLAN	Description	Edge VLAN Mode	Trust ARP Responses	Trust DHCP Responses
1		✓	✗	✗

Type to search in tables Row Count: 1

Add **Edit** **Delete** **Exit**

FIGURE 249 Profile Overrides - Network Bridge VLAN screen

4. Review the following VLAN configuration parameters to determine whether an override is warranted:

VLAN	Lists the numerical identifier defined for the Bridge VLAN when it was initially created. The available range is from 1 - 4095. This value cannot be modified during the edit process.
Description	Lists a description of the VLAN assigned when it was created or modified. The description should be unique to the VLAN's specific configuration and help differentiate it from other VLANs with similar configurations.
Edge VLAN Mode	Defines whether the VLAN is currently in edge VLAN mode. A green checkmark defines the VLAN as extended. An edge VLAN is the VLAN where hosts are connected. For example, if VLAN 10 is defined with wireless clients, and VLAN 20 is where the default gateway resides, VLAN 10 should be marked as an edge VLAN and VLAN 20 shouldn't be marked as an edge VLAN. When defining a VLAN as edge VLAN, the firewall enforces additional checks on hosts in that VLAN. For example, a host cannot move from an edge VLAN to another VLAN and still keep firewall flows active.
Trust ARP Response	When ARP trust is enabled, a green checkmark displays. When disabled, a red "X" displays. Trusted ARP packets are used to update the IP-MAC Table to prevent IP spoof and arp-cache poisoning attacks.
Trust DHCP Responses	When DHCP trust is enabled, a green checkmark displays. When disabled, a red "X" displays. When enabled, DHCP packets from a DHCP server are considered trusted and permissible. DHCP packets are used to update the DHCP Snoop Table to prevent IP spoof attacks.

5. Select **Add** to define a new Bridge VLAN configuration, **Edit** to modify or override an existing Bridge VLAN configuration or **Delete** to remove a VLAN configuration.

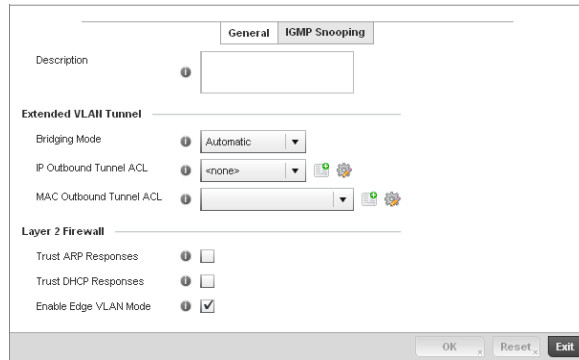


FIGURE 250 Bridge VLAN screen

6. The **General** tab displays by default.
7. If adding a new Bridge VLAN configuration, use the spinner control to define or override a **VLAN** ID between 1 - 4094. This value must be defined and saved before the General tab can become enabled and the remainder of the settings defined. VLAN IDs 0 and 4095 are reserved and unavailable.
8. Set or override the following General Bridge VLAN parameters:

Description	If creating a new Bridge VLAN, provide a description (up to 64 characters) unique to the VLAN's specific configuration to help differentiate it from other VLANs with similar configurations.
--------------------	---

9. Set or override the following **Extended VLAN Tunnel** parameters:

Bridging Mode	Specify one of the following bridging mode for use on the VLAN. <ul style="list-style-type: none"> • <i>Automatic</i>: Select automatic mode to let the controller determine the best bridging mode for the VLAN. • <i>Local</i>: Select Local to use local bridging mode for bridging traffic on the VLAN. • <i>Tunnel</i>: Select Tunnel to use a shared tunnel for bridging traffic on the VLAN. • <i>isolated-tunnel</i>: Select isolated-tunnel to use a dedicated tunnel for bridging traffic on the VLAN.
----------------------	--

IP Outbound Tunnel ACL	Select an IP Outbound Tunnel ACL for outbound traffic from the pulldown menu. If an appropriate outbound IP ACL is not available click the create button to make a new one.
-------------------------------	---

MAC Outbound Tunnel ACL	Select a MAC Outbound Tunnel ACL for outbound traffic from the pulldown menu. If an appropriate outbound MAC ACL is not available click the create button to make a new one.
--------------------------------	--

NOTE

Local and Automatic bridging modes do not work with ACLs. ACLs can only be used with tunnel or isolated-tunnel modes.

10. Set or override the following **Layer 2 Firewall** parameters:

- Trust ARP Response** Select the checkbox to use trusted ARP packets to update the DHCP Snoop Table to prevent IP spoof and arp-cache poisoning attacks. This feature is disabled by default.
- Trust DHCP Responses** Select the checkbox to use DHCP packets from a DHCP server as trusted and permissible within the managed network. DHCP packets are used to update the DHCP Snoop Table to prevent IP spoof attacks. This feature is disabled by default.
- Edge VLAN Mode** Select the checkbox to enable edge VLAN mode. When selected, the edge controller's IP address in the VLAN is not used for normal operations, as its now designated to isolate devices and prevent connectivity. This feature is enabled by default.

11. Select the **OK** button to save the changes to the General tab. Select **Reset** to revert to the last saved configuration.

12. Select the **IGMP Snooping** tab to define the VLAN's IGMP configuration.

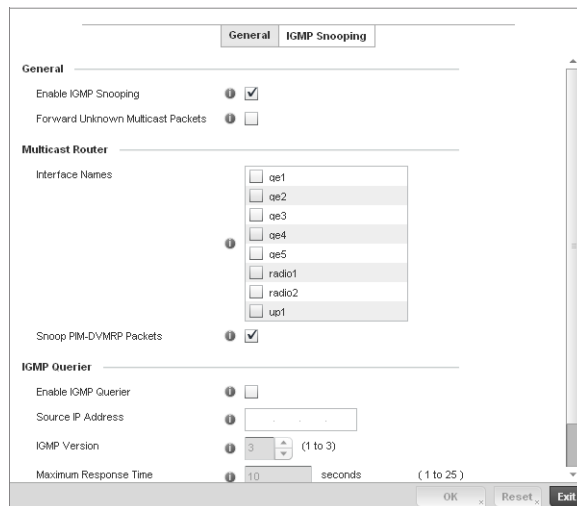


FIGURE 251 Bridge VLAN screen - IGMP Snooping Tab

13. Define the following **IGMP Snooping** parameters for the Bridge VLAN configuration:

The *Internet Group Management Protocol* (IGMP) is a protocol used for managing members of IP multicast groups. The controller listens to IGMP network traffic and forwards the IGMP multicast packets to radios on which the interested hosts are connected. On the wired side of the network, the controller floods all the wired interfaces. This feature reduces the unnecessary flooding of multicast traffic in the network.

- Enable IGMP Snooping** Select the check box to enable IGMP snooping on the controller. If disabled, snooping on a per VLAN basis is also disabled. This feature is enabled by default. If disabled, the settings under bridge configuration are overridden. For example, if IGMP snooping is disabled, but the bridge VLAN is enabled, the effective setting is disabled.
- Forward Unknown Unicast Packets** Select the check box to enable the controller to forward multicast packets from unregistered multicast groups. If disabled (the default setting), the unknown multicast forward feature is also disabled for individual VLANs.

14. Within the **Multicast Router** section, check the boxes of those interfaces used by the controller as a multicast router interface. Multiple controller interfaces can be selected and overridden.
15. Optionally select the **Snoop PIM-DVMRP Packets** box to snoop packets across the selected interface(s). This option is enabled by default.
16. Set the following **IGMP Querier** parameters for the profile's bridge VLAN configuration:

Enable IGMP Querier	Select the check box to enable IGMP querier. IGMP snoop querier is used to keep host memberships alive. It's primarily used in a network where there's a multicast streaming server and hosts subscribed to the server and no IGMP querier present. The controller can perform the IGMP querier role. An IGMP querier sends out periodic IGMP query packets. Interested hosts reply with an IGMP report packet. IGMP snooping is only conducted on wireless radios. IGMP multicast packets are flooded on wired ports. IGMP multicast packets are not flooded on the wired port. IGMP membership is also learnt on it and only if present, then it is forwarded on that port. A Brocade Mobility AP-7131 model access point can also be an IGMP querier.
Source IP Address	Define an IP address applied as the source address in the IGMP query packet. This address is used as the default VLAN querier IP address.
IGMP Version	Use the spinner control to set the IGMP version compatibility to either version 1, 2 or 3. The default setting is 3.
Maximum Response Time	Specify the maximum time (between 1 - 25 seconds) before sending a responding report. When no reports are received from a radio, radio information is removed from the snooping table. The controller only forwards multicast packets to radios present in the snooping table. For IGMP reports from wired ports, the controller forwards these reports to the multicast router ports. The default setting is 10 seconds.
Other Querier Timer Expiry	Specify an interval in either <i>Seconds</i> (60 - 300) or <i>Minutes</i> (1 - 5) used as a timeout interval for other querier resources. The default setting is 1 minute.

17. Select the **OK** button located at the bottom right of the screen to save the changes to the IGMP Snooping tab. Select **Reset** to revert to the last saved configuration.

Cisco Discovery Protocol Configuration

[Profile Network Configuration](#)

The Cisco Discovery Protocol (CDP) is a proprietary Data Link Layer network protocol implemented in Cisco networking equipment and used to share information about network devices.

To override *Cisco Discovery Protocol* (CDP) configuration:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.
2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.
3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Network** to expand its sub menu options.

5. Select Cisco Discovery Protocol.

FIGURE 252 Profile Overrides - Network Cisco Discovery Protocol screen

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

6. Check the **Enable CDP** box to enable Cisco Discovery Protocol on the device.
7. Refer to the **Hold Time** field and use the spinner control to define a hold time between 10 - 1800 seconds for transmitted CDP Packets. The default value is 180 seconds.
8. Refer to the **Timer** field and use the spinner control to define a interval between 5 - 900 seconds to transmit CDP Packets. The default value is 60 seconds.
9. Select the **OK** button to save the changes and overrides. Select **Reset** to revert to the last saved configuration.

Link Layer Discovery Protocol Configuration

Profile Network Configuration

The Link Layer Discovery Protocol (LLDP) or IEEE 802.1AB is a vendor-neutral Data Link Layer protocol used by network devices for advertising of (announcing) their identity, capabilities, and interconnections on a IEEE 802 LAN network. The protocol is formally referred to by the IEEE as Station and Media Access Control Connectivity Discovery. Both LLDP snooping and ability to generate and transmit LLDP packets will be provided.

Information obtained via CDP and LLDP snooping is available in the UI. In addition, information obtained via CDP / LLDP snooping is provided by an AP during the adoption process, so the L2 switch device name detected by the AP can be used as a criteria in the provisioning policy.

To override *Link Layer Discovery Protocol (LLDP)* configuration:

1. Select **Devices** from the Configuration tab.

The Device Configuration screen displays a list of managed devices or peer controllers. The listed devices can either be other controllers or Access Points within the managed network.

2. Select a target device (by double-clicking it) from amongst those displayed within the Device Configuration screen.

Devices can also be selected directly from the Device Browser in the lower, left-hand, side of the UI.

3. Select **Profile Overrides** from the Device menu to expand it into sub menu options.
4. Select **Network** to expand its sub menu options.
5. Select **Link Layer Discovery Protocol**.

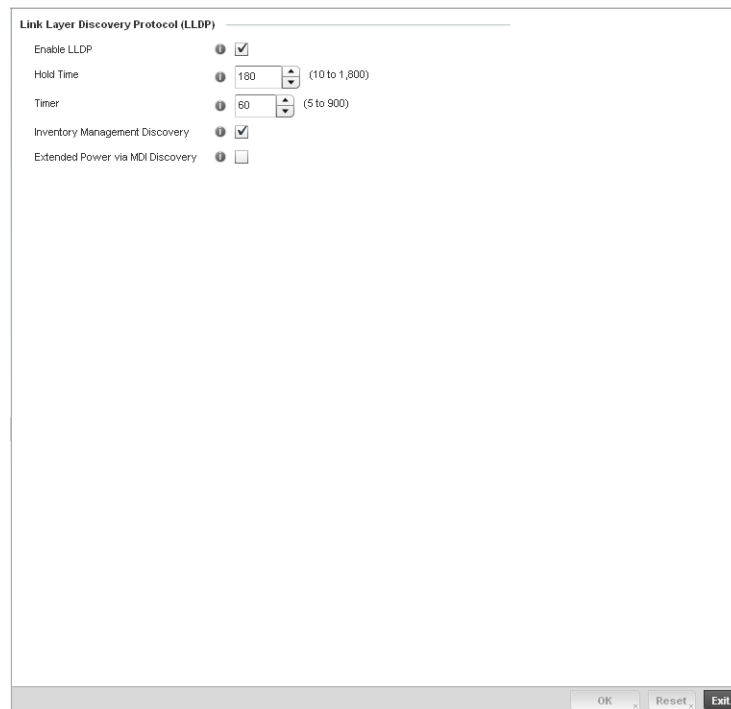


FIGURE 253 Profile Overrides - Network Link Layer Discovery Protocol screen

NOTE

A blue override icon (to the left of a parameter) defines the parameter as having an override applied. To remove an override go to the **Basic Configuration** section of the device and click the **Clear Overrides** button. This will remove all overrides from the device.

6. Check the **Enable LLDP** box to enable Link Layer Discovery Protocol on the device.
7. Refer to the **Hold Time** field and use the spinner control to define a hold time between 10 - 1800 seconds for transmitted LLDP Packets. The default value is 180 seconds.

8. Refer to the **Timer** field and use the spinner control to define the interval between 5 - 900 seconds to transmit LLDP Packets. The default value is 60 seconds.
9. Check the **Inventory Management Discovery** box to enable this feature. Inventory Management Discovery is used to track and identify inventory attributes including manufacturer, model, or software version.
10. Select the **Extended Power via MDI Discovery** box to enable this feature. Extended Power via MDI Discovery provides detailed power information through end points and other connected devices.
11. Select the **OK** button to save the changes and overrides. Select **Reset** to revert to the last saved configuration.

Miscellaneous Network Configuration

Profile Network Configuration

A controller profile can be configured to include a hostname in a DHCP lease for a requesting device and its profile. This helps an administrator track the leased DHCP IP address by hostname for the controller supported device profile. When numerous DHCP leases are assigned, an administrator can better track the leases when hostnames are used instead of devices.

To include a hostnames in DHCP request:

1. Select **Configuration > Profiles > Network**.
2. Expand the Network menu to display its submenu options
3. Select **Miscellaneous**.

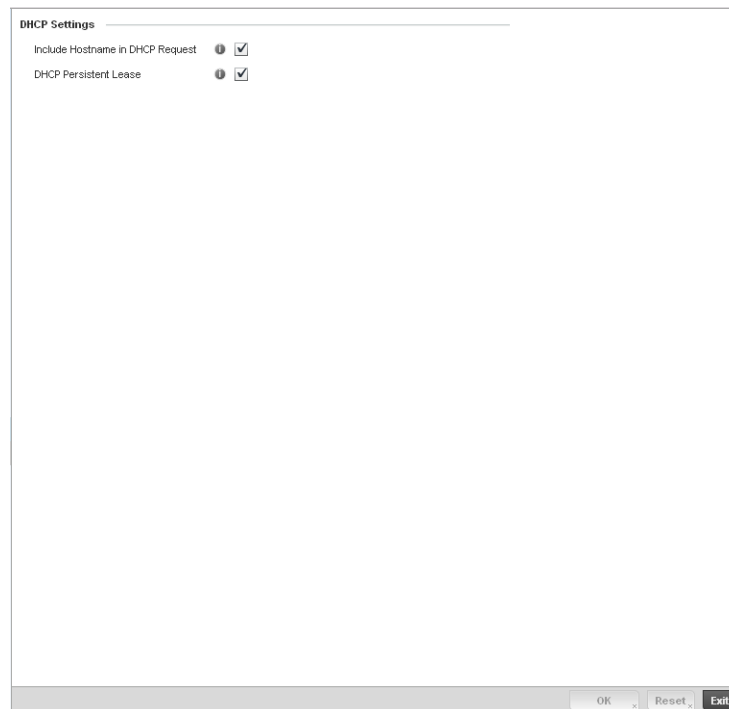


FIGURE 254 Profile Miscellaneous screen

4. Refer to the DHCP Settings section to configure miscellaneous DHCP Settings.

Include Hostname in DHCP Request

Select the Include Hostname in DHCP Request checkbox to include a hostname in a DHCP lease for a requesting device. This feature is disabled by default.

DHCP Persisten Lease

Check this box to enable a persistent DHCP lease for the device. A persistent DHCP lease assigns the same IP Address and other network information to the device each time it renews its DHCP lease.

5. Select the **OK** button located at the bottom right of the screen to save the changes. Select **Reset** to revert to the last saved configuration.

Profile Network Configuration and Deployment Considerations

Profile Network Configuration

Before defining a profile's network configuration, refer to the following deployment guidelines to ensure the profile configuration is optimally effective:

- Administrators often need to route traffic to interoperate between different VLANs. Bridging VLANs are only for non-routable traffic, like tagged VLAN frames destined to some other device which will untag it. When a data frame is received on a port, the VLAN bridge determines the associated VLAN based on the port of reception.
- Static routes while easy can be overwhelming within a large or complicated network. Each time there is a change, someone must manually make changes to reflect the new route. If a link goes down, even if there is a second path, the router would ignore it and consider the link down.
- Static routes require extensive planning and have a high management overhead. The more routers that exist in a network, the more routes needing to be configured. If you have N number of routers and a route between each router is needed, then you must configure $N \times N$ routes. Thus, for a network with nine routers, you'll need a minimum of 81 routes ($9 \times 9 = 81$).

Profile Security Configuration

A controller or Access Point profile can have its own firewall policy, wireless client role policy, WEP shared key authentication, NAT policy and VPN policy (controller only) applied. If an existing firewall, client role or NAT policy is unavailable, an administrator can be navigated from the **Configuration > Profiles** section of the controller UI to the **Configuration > Security** portion of the UI to create the required security policy configuration. Once created, separate policies can be applied to the controller profile to best support the data protection and security requirements in respect to the controller or Access Point model being supported by the profile.

For more information, refer to the following sections:

- [Defining Security Settings](#)
- [Setting the Certificate Revocation List \(CRL\) Configuration](#)
- [Configuring ISAKMP Policies](#)
- [Configuring Transform Sets](#)
- [Setting the Profile's VPN Configuration](#)
- [Setting the Profile's NAT Configuration](#)

Defining Security Settings

Profile Security Configuration

A controller profile can leverage existing firewall, wireless client role and WIPS policies and configurations and apply them to the profile's configuration. This affords each controller profile a truly unique combination of data protection policies best meeting the data protection requirements of that controller profile.

To define a profile's security settings:

1. Select the Configuration tab from the Web UI.
2. Select **Profiles** from the Configuration tab.
3. Select **Manage Profiles** from the Configuration > Profiles menu.
4. Select **Security**.
5. Select **Settings**.

The screenshot shows a web-based configuration interface for security settings. It is divided into two main sections: 'General' and 'Wireless IDS/IPS'. In the 'General' section, there are three configuration items: 'Firewall Policy' (dropdown menu set to '<none>'), 'Wireless Client Role Policy' (dropdown menu set to '<none>'), and 'WEP Shared Key Authentication' (checkbox, currently unchecked). The 'Wireless IDS/IPS' section contains one item: 'Advanced WIPS Policy' (dropdown menu). At the bottom right of the interface, there are three buttons: 'OK', 'Reset', and 'Exit'.

FIGURE 255 Security - Settings screen

6. Refer to the **General** field to assign or create the following security policy's to the profile:

Firewall Policy

Use the drop-down menu to select an existing Firewall Policy to use as an additional security mechanism with this controller profile. All devices using this controller profile and Access Point must meet the requirements of the firewall policy to access the managed network. A firewall is a mechanism enforcing access control, and is considered a first line of defense in protecting proprietary information within the wireless controller managed network. The means by which this is accomplished varies, but in principle, a firewall can be thought of as mechanisms both blocking and permitting data traffic within the wireless controller managed network. If an existing Firewall policy does not meet your requirements, select the **Create** icon to create a new firewall policy that can be applied to this profile. An existing policy can also be selected and edited as needed using the **Edit** icon. For more information, see [Wireless Firewall on page 9-419](#) and [Configuring a Firewall Policy on page 9-420](#).

Wireless Client Role Policy

Use the drop-down menu to select a client role policy the controller uses to strategically filter client connections based on a pre-defined set of filter rules and connection criteria. If an existing Wireless Client Role policy does not meet your requirements, select the **Create** icon to create a new configuration that can be applied to this profile. An existing policy can also be selected and edited as needed using the **Edit** icon. For more information, see [Wireless Client Roles on page 9-438](#).

WEP Shared Key Authentication

Select the check box to require devices using this profile to use a WEP key to access the managed network using this profile. The wireless controller, other proprietary routers, and Brocade Mobility Solutions clients use the key algorithm to convert an ASCII string to the same hexadecimal number. Clients without Brocade Mobility Solutions adapters need to use WEP keys manually configured as hexadecimal numbers. This option is disabled by default.

NOTE

Advanced WIPS Policy is only supported on wireless controllers and requires a dedicated WIPS sensor, but does not require a sensor license. Standard WIPS is available on all RF Domain managers and supports on channel, off channel and dedicated sensor scanning modes.

7. Select an **Advanced WIPS Policy** from the drop-down menu. Define an advanced WIPS configuration to optionally remove (terminate) unwanted device connections, and sanction (allow) or unsanction (disallow) specific events within the managed network.

If an existing Advanced WIPS policy does not meet the profile's data protection requirements, select the **Create** icon to create a new configuration that can be applied to the profile. An existing policy can also be selected and edited as needed using the **Edit** icon. For more information, see [Configuring an Advanced WIPS Policy on page 9-456](#).

8. Select **OK** to save the changes made within the Settings screen. Select **Reset** to revert to the last saved configuration.

Setting the Certificate Revocation List (CRL) Configuration

[Profile Security Configuration](#)

A *certificate revocation list* (CRL) is a list of certificates that have been revoked or are no longer valid. A certificate can be revoked if the *certificate authority* (CA) had improperly issued a certificate, or if a private-key is compromised. The most common reason for revocation is the user no longer being in sole possession of the private key.

To define a CRL configuration that can be applied to a controller profile:

1. Select the Configuration tab from the Web UI.
2. Select **Profiles** from the Configuration tab.
3. Select **Manage Profiles** from the Configuration > Profiles menu.
4. Select **Security**.
5. Select **Certificate Revocation**.

Certificate Revocation List (CRL) Update Interval

Trustpoint Name	URL	Hours	
TEST1	MOTO.COM	5 hour	

OK Reset Exit

FIGURE 256 Security - Certificate Revocation screen

6. Select the **+ Add Row** button to add a column within the **Certificate Revocation List (CRL) Update Interval** table to quarantine certificates from use in the managed network.

Additionally, a certificate can be placed on hold for a defined period. If, for instance, a private key was found and nobody had access to it, its status could be reinstated.

 - b. Provide the name of the trustpoint in question within the **Trustpoint Name** field. The name cannot exceed 32 characters.
 - c. Enter the resource ensuring the trustpoint's legitimacy within the **URL** field.
 - d. Use the spinner control to specify an interval (in hours) after which a device copies a CRL file from an external server and associates it with a trustpoint.
7. Select **OK** to save the changes made within the Certificate Revocation screen. Select **Reset** to revert to the last saved configuration.

Configuring ISAKMP Policies

Profile Security Configuration

ISAKMP (also known as IKE) is the negotiation protocol enabling two hosts to agree on how to build an IPsec security association. To configure the security appliance for virtual private networks, set global parameters that apply system wide and define policies peers negotiate to establish a VPN tunnel.

The ISAKMP protocol is an IPsec standard protocol used to ensure security for VPN negotiation, and remote host or network access. ISAKMP provides an automatic means of negotiation and authentication for communication between two or more parties. ISAKMP manages IPsec keys automatically.

To configure ISAKMP on the wireless controller:

1. Select the Configuration tab from the Web UI.
2. Select **Profiles** from the Configuration tab.
3. Select **Manage Profiles** from the Configuration > Profiles menu.
4. Select **Security**.
5. Select **ISAKMP Policy**.

The ISAKMP screen displays by default. The ISAKMP screen lists those ISAKMP policies created thus far. Use the *ISAKMP Policy* screen to configure the *Internet Security Association and Key Management Protocol (ISAKMP)* for creating a VPN. ISAKMP is a framework for authentication and key exchange. It defines the procedures and packet formats to establish, negotiate, modify and delete *Security Associations (SA)*. Any of these policies can be selected and applied to the controller.

A VPN tunnel is negotiated in two phases. The first phase creates an ISAKMP SA that's a control channel. The data channel is negotiated using this control channel. ISAKMP policy parameters are not negotiated and the transform set is for negotiating the data channel (IPsec SAs).

The screenshot shows a configuration window titled "Settings" for an ISAKMP Policy. It includes the following fields and options:

- Authentication Type:** A dropdown menu set to "Pre-Shared".
- Encryption:** A dropdown menu set to "AES-256".
- DH Group Identifier:** A text input field containing "2", with "(1-2,5)" shown as a range.
- Hash Algorithm:** Radio buttons for "MD5" and "SHA", with "SHA" selected.
- SA Lifetime:** A text input field containing "1", followed by a "Days" dropdown menu and "(0 to 24,855)" shown as a range.

At the bottom right of the window are three buttons: "OK", "Reset", and "Exit".

FIGURE 258 ISAKMP Policy screen

8. Set the following configuration parameters for the new or modified ISAKMP policy:

ISAKMP Policy

If creating a new ISAKMP policy, assign it name to help differentiate it from others that may have a similar configuration. The policy name cannot exceed 64 characters. The name cannot be modified as part of the edit process.

Authentication Type

Set the key sharing mechanism used for establishing a secure connection between two peers using this ISAKMP policy. Use the drop-down menu to select one the following two authentication options:

Pre-Shared – Select this option to use a key shared between the VPN endpoints. Pre-Shared is the default setting.

RSA Signature – Uses a RSA signature as the authentication key. Ensure digital certificates and RSA keys have been installed on the target system before using this option.

Encryption

Use the drop-down menu to select the encryption algorithm. Select from the following options:

DES – DES stands for *Data Encryption Standard*. It uses a 56-bit key for encryption. This standard is deprecated and replaced by the 3DES standard. It is provided for backward compatibility.

3DES - 3DES or *Triple DES* is a encryption standard that replaced DES. It provides a simple method of increasing the key size of the DES algorithm to protect from brute force attacks. It uses a set of three (3) standard 56-bit DES keys to provide increased key length for encryption.

AES – AES stands for *Advanced Encryption Standard* or the Rijndael Encryption Algorithm that was adopted as the new FIPS standard in the year 2002. It is a symmetric-key encryption standard that uses three (3) block ciphers of length 128, 192, 256 bits. This option represents the AES-128 bit block cipher.

AES-192 – This option represents the AES-192 bit block cipher.

AES-256 – This option represents the AES-256 bit block cipher. AES -256 is the default setting.

DH Group Identifier	<p>Set the <i>Diffie-Hellman</i> (DH) group identifier used by this ISAKMP policy. DH is a cryptographic protocol that allows 2 entities that have no prior knowledge of each other to jointly derive and establish a shared secret key over an unsecured communication channel. This secret key can then be used to initiate a secure connection between the two entities.</p> <p>The valid values are 1, 2 and 5 and indicates the group used for the key exchange. The default setting is 2.</p>
Hash Algorithm	<p>Set the hash algorithm. Select from:</p> <p><i>MD5</i> – MD5 or <i>Message-Digest algorithm 5</i> is a popular 128-bit hash-function. It is commonly used for checking the integrity of files.</p> <p><i>SHA</i> – <i>Secure Hash Algorithm</i> (SHA) is a NIST certified FIPS hash algorithm. SHA is the default setting.</p>
SA Lifetime	<p>Set the lifetime in seconds for the <i>security association</i> (SA) used by this ISAKMP policy. This is the lifetime of the ISAKMP SA. The lifetime for ESP/AH SAs are configured separately.</p> <p><i>Days</i> – Sets the SA duration in days (1 - 24,856).</p> <p><i>Hours</i> – Sets the SA duration in hours (1 - 596,524).</p> <p><i>Minutes</i> – Sets the SA duration in minutes (1 - 35, 791, 395).</p> <p><i>Seconds</i> – Sets the SA duration in seconds (60 - 2,147,483,646). The default setting is 86,400 seconds.</p>

9. Select **OK** to save the ISAKMP policy configuration. Select **Reset** to revert to the last saved configuration.

Configuring Transform Sets

[Profile Security Configuration](#)

Use the **Transform Set** screen to configure and manage transform sets. A Transform Set is a set of parameters that transform an IP packet from clear text to cipher text. The transform set is an acceptable combination of security protocols, algorithms and other settings that are applied to IPSec protected traffic.

With manually established security associations, there's no negotiation with the peer. Both sides must specify the same transform set, regardless of whether the SA is manual or automatic. For manual SAs, the ISAKMP policy does not apply. If you change a Transform Set definition, the change is only applied to Crypto Map entries that reference the Transform Set. If a transform-set is changed, the existing SAs are removed.

To define a transform set configuration that can be applied to a controller profile:

1. Select the Configuration tab from the Web UI.
2. Select **Profiles** from the Configuration tab.
3. Select **Manage Profiles** from the Configuration > Profiles menu.
4. Select **Security**.
5. Select **Transform Set**.

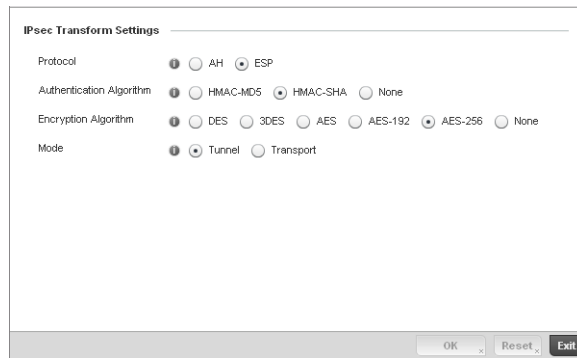


FIGURE 260 Transform Set Configuration Add screen

8. Set the following configuration parameters for the new or modified Transform Set:

Transform Set	If creating a new Transform Set, assign it a name to help differentiate it from others that may have a similar configuration. The name cannot exceed 64 characters. The name cannot be modified as part of the edit process.
Protocol	Select the check box of the IPsec protocol used with the Transform Set. AH provides data authentication only. ESP provides data confidentiality and well as authentication. ESP is the default setting.
Authentication Algorithm	Set the authentication algorithm used to validate identity. <i>HMAC-MD5</i> – Use the <i>Message Digest 5 (MD5)</i> as the HMAC algorithm. <i>HMAC-SHA</i> – Use the <i>Secure Hash Algorithm (SHA)</i> as the HMAC algorithm. HMAC-SHA is the default setting. <i>None</i> – Applies no authentication. If the protocol is AH, None cannot be selected.
Encryption Algorithm	The encryption algorithm check box only displays when ESP is selected as the Transform Set protocol. By default, the Transform set uses AES-256. AES stands for Advanced Encryption Standard. It's a symmetric-key encryption standard that uses a block ciphers length 256 bits. AES -256 is the default setting. Selecting <i>None</i> applies no encryption. When the protocol is ESP, encryption and authentication cannot both be set to None.
Mode	Set the mode used for packet organization with respect to header location and the scope of ESP or AH protection boundary. Use Tunnel for site-to-site VPN and Transport mode for remote VPN configurations. The default mode is Tunnel.

9. Select **OK** to save the Transform Set policy configuration. Select **Reset** to revert to the last saved configuration.

Setting the Profile's VPN Configuration

Profile Security Configuration

A *Virtual Private Network (VPN)* provides secure communications between two or more IP networks or computer devices across an IP network infrastructure. VPN is commonly deployed to securely inter-connect remote sites or provide remote access across the public Internet and connect networks across a private IP network. Traffic exchanged between VPN end-points can transparently pass through shared resources such as routers, switches, and other telecommunications equipment and is secured using *IP Security (IPSec)*.

To define a VPN configuration that can be applied to a controller profile:

1. Select the Configuration tab from the Web UI.
2. Select **Profiles** from the Configuration tab.
3. Select **Manage Profiles** from the Configuration > Profiles menu.
4. Select **Security**.
5. Select **VPN Configuration**.
6. The **Basic Settings** tab displays by default. Refer to the Peer Settings table to add peer addresses and keys for VPN tunnel destinations. Use the **+ Add Row** function as needed to add additional destinations and keys.

The screenshot displays the 'Basic Settings' tab for VPN configuration. It features a 'Peer Settings' table with columns for 'Peer Address', 'Key', and a delete icon. Below this is an 'Additional Settings' section with expandable options for 'ISAKMP Keepalive', 'IPsec Lifetime (seconds)', and 'IPsec Lifetime (kB)'. The 'Aggressive Mode Peer (ISAKMP Responder)' section contains a table with columns for 'Peer Type', 'Peer Id', 'Preshared Key', and a delete icon. The interface includes 'Add Row' buttons for both tables and 'OK', 'Reset', and 'Exit' buttons at the bottom.

FIGURE 261 Security VPN Configuration - Basic Settings screen

7. Set the **Peer Address** from (within the **Peer Settings** table) for a known device at the other end of the managed VPN tunnel.
8. Define a **Key** (between 8 - 21 characters long) shared between the peer and a device (and is required by devices at both ends of the tunnel for inter operation). The key is exposed as a series of “*****” characters unless the **Show** radio button is selected to display the actual characters comprising the key. This is the key set for the policy.
9. Select the **More >** button to the right of the **Additional Settings** field to expand the screen to display several IPsec parameters and the **Aggressive Mode Peer** table.

10. Set the following IPsec parameters within the **Additional Settings** field:

IPsec ISAKMP Keep Alive	Use the spinner control to set the IPsec ISAKMP keep alive duration (between 10 - 3,600) seconds. This determines the frequency of IKE keep alive messages for dead peer detection. The default is 10 seconds.
IPsec Lifetime (seconds)	Set the IPsec Lifetime value in either <i>Seconds</i> (90 - 2,147,483,646), <i>Minutes</i> (2 - 35,791,395), <i>Hours</i> (1 - 596,524) or <i>Days</i> (1 - 24,856). This value represents the time-based IPsec security association lifetime. The default setting is 3,600 seconds. The IPsec Lifetime this is a global setting applying to all VPN Security associations. This global lifetime value can be overridden by crypto maps.
IPsec Lifetime (kB)	Sets the IPsec Lifetime value in kilo bytes (500 - 2,147,483, 646). The life time of the VPN tunnel is decided by the amount of traffic that passes through the VPN. Once this value is exceeded, the VPN is automatically closed and the clients have to re-negotiate the security association. The default is 4,608,000 kilobytes. The IPsec Lifetime this is a global setting applying to all VPN Security associations. This global lifetime value can be overridden by crypto maps.

NOTE

The security association is re-negotiated when the IPsec Lifetime value expires, whether it's seconds or kB, whichever comes first.

11. Refer to the **Aggressive Mode Peer (ISAKMP Responder)** table to set the following parameters:

Aggressive mode is an IKE policy mode which reduces the number of steps during IKE authentication negotiation between two IKE peers. This mode has less negotiation power and does not provide identity protection between the peers. This mode is vulnerable to man-in-the-middle attacks.

Peer Type	Define the type of aggressive mode peer used as one of the following. <i>IP Address</i> - The IP address of the peer is provided. <i>Hostname</i> - The Host Name of the peer is provided. <i>Distinguished Name</i> - The peer type is defined by a distinguished name.
Peer ID	Set the Peer ID as either an IP address, hostname or distinguished name.
Preshared Key	Provide a shared key between this device and the aggressive mode peer. Both use the same key to interoperate. The key is exposed as a series of "*****" characters unless the Show check box is selected to display the actual characters comprising the key.

12. Select the **+ Add Row** button to add additional rows.

13. Select **OK** to update the VPN policy's Basic Settings. Select **Reset** to revert to the last saved configuration.

14. Select the **Crypto Map** tab.




FIGURE 262 Security VPN Configuration - Crypto Map screen

Crypto Map entries are sets of configuration parameters for encrypting packets that pass through the VPN Tunnel. The screen lists the various Crypto Map entries defined under a particular VPN Policy.

15. Either select **Add** to create a new Crypto Map configuration, **Edit** to modify the attributes of an existing Crypto Map configuration or **Delete** to permanently remove a selected configuration.
16. Provide a **Name** for the Crypto Map up to 64 characters in length. The name should help distinguish the Crypto Map from others with similar configurations.
17. Select **Add** to display a **Crypto Map Entry** screen where the attributes of the Crypto Map configuration can be defined.

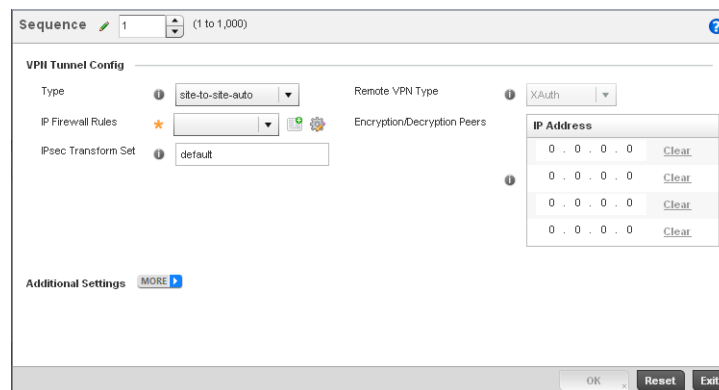


FIGURE 263 Crypto Map Entry screen

18. Use the spinner control to set a **Sequence** between 1 - 1000 that identifies the Crypto Map in the list of Crypto Map entries.

19. Define the following **VPN Tunnel Config** parameters:

Type	Define the type of tunnel to be created. Select from one of the following: <i>site-to-site-auto</i> – Sets a VPN connection (using ISAKMP) to a remote site where the connection parameters are configured automatically. This is the default setting. <i>remote-vpn</i> – Creates a VPN connection to a remote client. <i>site-to-site-manual</i> – Creates a VPN connection to a remote site where the connection parameters are configured manually.
IP Firewall Rules	Set the Firewall rules applied to this VPN tunnel. Select existing rules from the drop down list. These firewall rules are used to classify traffic for VPN. Use the Create icon to create a new Firewall rule set or the Edit icon to modify the configuration of an existing Firewall rule. For more information, see Configuring IP Firewall Rules on page 9-431 .
IPSec Transform Set	Define the existing (previously configured) transform set for this VPN Tunnel.
Remote VPN Client Type	This parameter is available only when <i>Remote VPN</i> is selected in the Type field. Select from one of the following: <i>IPsec-L2TP</i> – The remote VPN client uses an IPsec Layer 2 Tunnelling Protocol. <i>XAuth</i> – The remote VPN client uses X-Windows Authorization.
Encryption/Decryption Peers	If <i>Automatic Site-to-Site VPN</i> or <i>Remote VPN Client</i> is selected as the Type, Enter up to 4 IP addresses for Encryption/Decryption Peers. For Auto VPN, multiple peers are configured for redundancy. An administrator is required to configure just one peer, unless redundancy is required and another peer is needed.

20. Select the **More >** link to expand the screen to add the following **Additional Settings** for the Crypto Map configuration:

Perfect Forward Secrecy (PFS)	Select the check box to enable PFS and set DH group or either 1, 2 or 5. Enable this option to set an authenticated key-agreement protocol that uses public key cryptography. <i>Perfect Forward Secrecy (PFS)</i> ensures a session key derived from public and private keys is not compromised if one of the private keys is compromised in the future. This feature is disabled by default.
Mode of Tunnel Initiation	Select the tunnel mode. Choose from: <i>Main</i> – The complete process of authentication and key exchange is followed when Main mode is selected. Main is the default value. <i>Aggressive</i> – A truncated process of authentication and key exchange is followed when Aggressive mode is selected.
Local Identity Type (ISAKMP Initiator)	Define how this end of the tunnel is identified. Select from: <i>IP Address</i> – This end of the tunnel is identified by an IP address. IP Address is the default setting. <i>Hostname (FQDN)</i> – This end of the tunnel is identified by a host name. <i>Distinguished Name</i> – This end of the tunnel is identified by a domain name.
Local Identity Value	Set the name of the host or domain (only applicable for the initiator). The name cannot exceed 64 characters. This option is not available when IP Address is the Local Identity Type.

21. Set the following **IPsec Security Association (SA)** parameters for the Crypto Map configuration:

- Lifetime (seconds)** Select the check box to enable the spinner control to set the Security Association lifetime in seconds (90 - 2,147,483,646). This feature is disabled by default.
- Lifetime (kb)** Select the check box to enable the spinner control to set the Security Association lifetime in kilobytes (500 - 2,147,483,646) for data traversing the VPN. This feature is disabled by default.
- Per Host SA** Select the check box to set the granularity for IPsec security associations to one per-host. This feature is disabled by default.

22. Select **OK** to save the Crypto Map policy configuration. Select **Reset** to revert to the last saved configuration.

23. Select the **Remote VPN** tab to define an additional remote VPN configuration beyond what has been defined in the Basic Settings tab.

FIGURE 264 Security VPN Configuration - Remote VPN screen

24. Use the **Local Pool Settings** table set a **Low Address** and **High Address** range for local IP address assignments from the local pool. This table represents the virtual IP pool that can be assigned to remote VPN clients. Select the **+ Add Row** button as required to additional entries in the table.

25. Refer to the **Remote VPN Config** field to set the following parameters:

- DNS Server** Enter the address of the DNS Server configured on the remote VPN client.
- WINS Server** Enter the address of the WINS Server configured on the remote VPN client.

26. Select the radio button corresponding to the desired **Authentication Method** used to authenticate the remote VPN client. Set to **RADIUS** to use a dedicated RADIUS Server (If using XAuth as the remote VPN Client Type) or **Local** for internal controller resources for authentication. Selecting **None** bypasses authenticating the remote client. Local is the default setting.
27. If using the Local default setting, select **+ Add Row** to add a user name for the local login account and the password for the user name on the local server. The password is exposed as a series of “*****” characters, unless the **Show** radio button is selected to display the actual characters comprising the password.
28. If selecting RADIUS, specify the **NAS** originating the RADIUS accept request and define a RADIUS configuration for remote client authentication within the RADIUS Configuration field.
29. If selecting RADIUS, specify an **ARP Timeout** value in *Days, Hours, Minutes* or *Seconds*.

RADIUS Configuration

NAS

RADIUS Servers

Server Type	Server IP	Authentication Port	Password	
				



FIGURE 265 Remote VPN Authentication screen - RADIUS Configuration

30. Set the following **RADIUS Configuration** parameters:

Server Type	Select the RADIUS server type. Select from: <i>Primary</i> – The selected server is the Primary RADIUS Server. <i>Secondary</i> – The selected server is the Secondary RADIUS Server.
Server IP	Set IP address of the RADIUS server.
Authentication Port	Select the check box and use the spinner control to set the UDP port number where the defines RADIUS server is listening.
Password	Enter the password to log on to the RADIUS server. The password is exposed as a series of “*****” characters unless the Show check box is selected to display the actual characters comprising the password.

31. Select **OK** to update the Remote VPN Authentication settings. Select **Reset** to revert to the last saved configuration.

Setting the Profile’s NAT Configuration

Profile Security Configuration

Network Address Translation (NAT) is a technique to modify network address information within IP packet headers in transit across a traffic routing device. This enables mapping one IP address to another to protect wireless controller managed network address credentials. With typical deployments, NAT is used as an IP masquerading technique to hide private IP addresses behind a single, public facing, IP address.

NAT is a process of modifying network address information in IP packet headers while in transit across a traffic routing device for the purpose of remapping one IP address to another. In most deployments NAT is used in conjunction with IP masquerading which hides RFC1918 private IP addresses behind a single public IP address.

NAT can provide a controller profile outbound Internet access to wired and wireless hosts connected to either a *thick* access point (such as an AP-7131, AP-7131N or AP-5131 model) or a RFS4000, RFS6000 or RFS7000 model wireless controller. Many-to-one NAT is the most common NAT technique for outbound Internet access. Many-to-one NAT allows a thick access point or wireless controller to translate one or more internal private IP addresses to a single, public facing, IP address assigned to a 10/100/1000 Ethernet port or 3G card.

To define a NAT configuration that can be applied to a controller profile:

1. Select the Configuration tab from the Web UI
2. Select **Profiles** from the Configuration tab.
3. Select **Manage Profiles** from the Configuration > Profiles menu.
4. Select **Security**.
5. Select **NAT**.




FIGURE 266 Security NAT screen

6. The **NAT Pool** displays by default. The NAT Pool screen lists those NAT policies created thus far. Any of these policies can be selected and applied to a controller profile.
7. Select **Add** to create a new NAT policy that can be applied to a controller profile. Select **Edit** to modify the attributes of an existing policy or select **Delete** to remove obsolete NAT policies from the list of those available to a controller profile.

Start IP	End IP	
192.31.98.135	192.31.98.150	

Add Row

OK Reset Exit

FIGURE 267 Security NAT Pool screen

8. If adding a new NAT policy or editing the configuration of an existing policy, define the following parameters:

Name	If adding a new NAT policy, provide a name to help distinguish it from others with similar configurations. The length cannot exceed 64 characters.
Prefix Length	Use the spinner control to set the netmask (between 1 - 30) of the network the pool address belongs to.
IP Address Range	Define a range of IP addresses that are hidden from the public Internet. NAT modifies network address information in the defined IP range while in transit across a traffic routing device. NAT only provides IP address translation and does not provide a firewall. A branch deployment with NAT by itself will not block traffic from being potentially routed through a NAT device. Consequently, NAT should be deployed with a stateful firewall.

9. Select the **+ Add Row** button as needed to append additional rows to the IP Address Range table.
10. Select **OK** to save the changes made to the profile's NAT Pool configuration. Select **Reset** to revert to the last saved configuration.
11. Select the **Static NAT** tab.
The Source tab displays by default.

The screenshot shows the 'Static NAT' configuration interface. At the top, there are three tabs: 'NAT Pool', 'Static NAT', and 'Dynamic NAT'. Below these are two more tabs: 'Source' and 'Destination'. The 'Source' tab is selected, and it contains a table with the following data:

Source IP	NAT IP	Network	
10.58.89.4	192.168.1.89	inside	

Below the table is an 'Add Row' button. At the bottom of the screen are 'OK', 'Reset', and 'Exit' buttons.

FIGURE 268 Static NAT screen

12. To map a source IP address from an internal network to a NAT IP address click the **+ Add Row** button. Enter the internal network IP address in **Source IP** field. Enter the NAT IP address in the **NAT IP** field.
13. Use the **Network** drop-down menu to set the NAT type either *Inside* or *Outside*. Select **Inside** to create a permanent, one-to-one mapping between an address on an internal network and a perimeter or external network. To share a Web server on a perimeter interface with the Internet, use static address translation to map the actual address to a registered IP address. Static address translation hides the actual address of the server from users on insecure interfaces. Casual access by unauthorized users becomes much more difficult. Static NAT requires a dedicated address on the outside network for each host. Inside NAT is the default setting.
14. Select the **Destination** tab to view destination NAT configurations and define packets passing through the NAT on the way back to the managed LAN are searched against to the records kept by the NAT engine. The destination IP address is changed back to the specific internal private class IP address to reach the LAN over the managed network.

Static NAT creates a permanent, one-to-one mapping between an address on an internal network and an external network. To share a Web server on an interface with the Internet, use static address translation to map the actual address to a registered IP address. Static address translation hides the actual server address from users on insecure interfaces. Casual access by unauthorized users becomes much more difficult. Static NAT requires a dedicated address on the outside network for each host.

Protocol	Select the protocol for use with static translation. TCP, UDP and Any are available options. TCP is a transport layer protocol used by applications requiring guaranteed delivery. It's a sliding window protocol handling both timeouts and retransmissions. TCP establishes a full duplex virtual connection between two endpoints. Each endpoint is defined by an IP address and a TCP port number. The <i>User Datagram Protocol</i> (UDP) offers only a minimal transport service, non-guaranteed datagram delivery, and provides applications direct access to the datagram service of the IP layer. UDP is used by applications not requiring the level of service of TCP or are using communications services (multicast or broadcast delivery) not available from TCP. The default setting is Any.
Destination IP	Enter the local address used at the (source) end of the static NAT configuration. This address (once translated) will not be exposed to the outside world when the translation address is used to interact with the remote destination.
Destination Port	Use the spinner control to set the local port used at the (source) end of the static NAT configuration. The default port is 1.
NAT IP	Enter the IP address of the matching packet to the specified value. The IP address modified can be either source or destination based on the direction specified.
NAT Port	Enter the port number of the matching packet to the specified value. This option is valid only if the direction specified is destination.
Network	Select Inside or Outside NAT as the network direction. Inside is the default setting.

17. Select **OK** to save the changes made to the static NAT configuration. Select **Reset** to revert to the last saved configuration.
18. Select the **Dynamic NAT** tab.

Dynamic NAT configurations translate the IP address of packets going out from one interface to another interface based on configured conditions. Dynamic NAT requires packets be switched through a NAT router to generate translations in the controller translation table.




FIGURE 272 Source ACL List screen

21. Set the following to define the Dynamic NAT configuration:

Source List ACL	Use the drop-down menu to select an ACL name to define the packet selection criteria for NAT. NAT is applied only on packets which match a rule defined in the access-list. These addresses (once translated) will not be exposed to the outside world when the translation address is used to interact with the remote destination.
Network	Select <i>Inside</i> or <i>Outside</i> NAT as the network direction for the dynamic NAT configuration. Inside is the default setting.
Interface	Use the drop-down menu to select the VLAN (between 1 - 4094) used as the communication medium between the source and destination points within the NAT configuration. Ensure the VLAN selected represents the intended network traffic within the NAT supported configuration. VLAN1 is available by default.
Overload Type	Select the check box of Overload Type used with the listed IP ACL rule. Options include <i>NAT Pool</i> , <i>One Global Address</i> and <i>Interface IP Address</i> . Interface IP Address is the default setting.
NAT Pool	Provide the name of an existing NAT pool for use with the dynamic NAT configuration.
Overload IP	Enables the use of one global address for numerous local addresses.

22. Select **OK** to save the changes made to the dynamic NAT configuration. Select **Reset** to revert to the last saved configuration.

Profile Security Configuration and Deployment Considerations

Profile Security Configuration

Before defining a profile's security configuration, refer to the following deployment guidelines to ensure the profile configuration is optimally effective:

- Make sure the contents of the Certificate Revocation List are periodically audited to ensure revoked certificates remained quarantined or validated certificates are reinstated.
- A RFS4000 model wireless controller ships with a baseline configuration supporting many-to-one NAT between devices connected to GE1 - GE5 ports on VLAN 1, and the UP1 port assigned to VLAN 2100. A RFS4000 can be deployed within a small site using its default configuration, and then be connected to a Internet service providing instant access to the Internet.

- NAT alone does not provide a firewall. If deploying NAT on a controller profile, add a firewall on the profile to block undesirable traffic from being routed. For outbound Internet access, a stateful firewall can be configured to deny all traffic. If port address translation is required, a stateful firewall should be configured to only permit the TCP or UDP ports being translated.
- A RFS6000 model wireless controller ships with a minimum baseline configuration without NAT enabled. A RFS6000 wireless controller requires VLAN configuration, IP addressing and NAT rules be created before many-to-one NAT services can be defined.
- Brocade Mobility RFS4000 and RFS6000 model wireless controllers can provide outbound NAT services for hosts connected to multiple VLANs. For small deployments, VLANs should be terminated within a RFS4000 wireless controller providing site routing services. For medium-scale deployments, VLANs are typically terminated on a L3 (IP layer) or L2 (Ethernet layer).

Profile Services Configuration

A controller profile can contain specific guest access (captive portal), DHCP server and RADIUS server configurations supported by the controller's own internal resources. These controller access, IP assignment and user authorization resources can be defined uniquely as controller profile requirements dictate.

To define a profile's services configuration:

1. Select the Configuration tab from the Web UI.
2. Select **Profiles** from the Configuration tab.
3. Select **Manage Profiles** from the Configuration > Profiles menu.
4. Select **Services**.

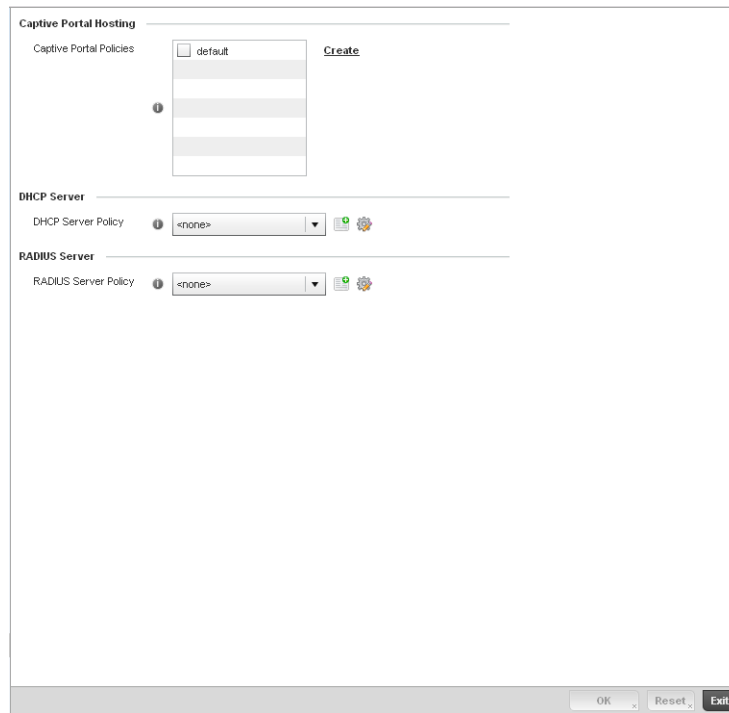


FIGURE 273 Profile Services screen

5. Refer to the **Captive Portal Hosting** section to select or set a controller guest access configuration (captive portal) for use with this profile.

A *captive portal* is guest access policy for providing guests temporary and restrictive access to the managed wireless network. The primary means of securing such controller guest access is a hotspot.

A captive portal policy's hotspot configuration provides secure authenticated controller access using a standard Web browser. Hotspots provides authenticated access by capturing and re-directing a wireless user's Web browser session to a captive portal login page where the user must enter valid credentials to access to the managed wireless network. Once logged into the managed hotspot, additional *Agreement*, *Welcome* and *Fail* pages provide the administrator with a number of options on the hotspot's screen flow and user appearance.

Either select an existing captive portal policy, use the default captive portal policy or select the **Create** link to create a new captive portal configuration that can be applied to this profile. For more information, see, [Configuring Captive Portal Policies on page 10-465](#).

6. Use the **DHCP Server Policy** drop-down menu assign this controller profile a DHCP server policy. If an existing DHCP policy does not meet the profile's requirements, select the **Create** button to create a new policy configuration that can be applied to this profile.

Dynamic Host Configuration Protocol (DHCP) allows hosts on an IP network to request and be assigned IP addresses as well as discover information about the managed network where they reside. Each subnet can be configured with its own address pool. Whenever a DHCP client requests an IP address, the DHCP server assigns an IP address from that subnet's address pool. When the controller's onboard DHCP server allocates an address for a DHCP client, the client is assigned a lease, which expires after an pre-determined interval. Before a lease expires, wireless clients (to which leases are assigned) are expected to renew them to continue

to use the addresses. Once the lease expires, the client is no longer permitted to use the leased IP address. The controller profile's DHCP server policy ensures all IP addresses are unique, and no IP address is assigned to a second client while the first client's assignment is valid (its lease has not expired).

Either select an existing captive portal policy or select the **Create** button to create a new captive portal configuration that can be applied to this profile. Existing policies can be modified by selecting the **Edit** icon. For more information, see [Setting the Controller's DHCP Configuration on page 10-477](#).

7. Use the **RADIUS Server Policy** drop-down menu to select an existing RADIUS server policy to use as a user validation security mechanism with this controller profile.

A controller profile can have its own unique RADIUS server policy to authenticate users and authorize access to the managed network. A profile's RADIUS policy provides the centralized management of controller authentication data (usernames and passwords). When a client attempts to associate to the controller, the controller sends the authentication request to the RADIUS server. An

AP-7131 model Access Point is also equipped with its own RADIUS server capability.

If an existing RADIUS server policy does not meet your requirements, select the **Create** button to create a new policy configuration that can be applied to this profile. Existing policies can be modified by selecting the **Edit** icon. For more information, see [Setting the Controller's RADIUS Configuration on page 10-491](#).

8. Select **OK** to save the changes made to the profile's services configuration. Select **Reset** to revert to the last saved configuration.

Profile Services Configuration and Deployment Considerations

[Profile Services Configuration](#)

Before defining a profile's captive portal, DHCP and RADIUS services configuration, refer to the following deployment guidelines to ensure the profile configuration is optimally effective:

- A profile plan should consider the number of wireless clients allowed on the profile's guest (captive portal) network and the services provided, or if the profile should support guest access at all.
- Controller profile configurations supporting a captive portal should include firewall policies to ensure logical separation is provided between guest and internal networks so internal networks and hosts are not reachable from managed guest devices.
- DHCP's lack of an authentication mechanism means a DHCP server supported controller profile cannot check if a client or user is authorized to use a given user class. This introduces a vulnerability when using user class options. Ensure a profile using the controller's internal DHCP resources is also provisioned with a strong user authorization and validation configuration.

Profile Management Configuration

The controller has mechanisms to allow/deny management access to the managed network for separate interfaces and protocols (*HTTP, HTTPS, Telnet, SSH* or *SNMP*). These management access configurations can be applied strategically to controller profiles as controller resource permissions dictate for the profile.

Additionally, an administrator can define a profile with unique configuration file and device firmware upgrade support. In a clustered environment, these operations can be performed on one controller, then propagated to each member of the cluster and onwards to the devices managed by each cluster member.

To define a profile's management configuration:

1. Select the Configuration tab from the Web UI.
2. Select **Profiles** from the Configuration tab.
3. Select **Manage Profiles** from the Configuration > Profiles menu.
4. Select **Management**.
5. Expand the Management menu item to display additional *Settings*, *Firmware* and *Heartbeat Management* options.
6. Select **Settings** from the Management menu.

FIGURE 274 Profile Management Settings screen

7. Refer to the **Management Policy** field to select or set a controller management configuration for use with this profile. A default management policy is also available if no existing policies are usable.

Use the drop-down menu to select an existing management policy to apply to this controller profile. If no management policies exist meeting the data access requirements of this controller profile, select the **Create** icon to access a series of screens used to define administration, access control and SNMP configurations. Select an existing policy and select the **Edit** icon to modify the configuration of an existing management policy. For more information, see [Viewing Management Access Policies on page 11-511](#).

8. Use to the **Critical Resource Policy** pulldown to set or override a critical resource policy for use with this profile. For more information on defining a critical resource policy, see *Managing Critical Resource Policies on page 5-223*.
9. Refer to the **Message Logging** field to define how the controller profile logs system events. It's important to log individual events to discern an overall pattern that may be negatively impacting controller performance using the configuration defined for this profile.

Enable Message Logging	Select the check box to enable the controller profile to log system events to a user defined log file or a syslog server. Selecting this check box enables the rest of the parameters required to define the profile's logging configuration. This option is disabled by default.
Remote Logging Host	Use this table to define numerical (non DNS) IP addresses for up to three external resources where logged system events can be sent on behalf of the controller profile. Select Clear as needed to remove an IP address.
Facility to Send Log Messages	Use the drop-down menu to specify the local server facility (if used) for the controller profile event log transfer.
Syslog Logging Level	Event severity coincides with the syslog logging level defined for the profile. Assign a numeric identifier to log events based on criticality. Severity levels include 0 - <i>Emergency</i> , 1 - <i>Alert</i> , 2 - <i>Critical</i> , 3 - <i>Errors</i> , 4 - <i>Warning</i> , 5 - <i>Notice</i> , 6 - <i>Info</i> and 7 - <i>Debug</i> . The default logging level is 4.
Console Logging Level	Event severity coincides with the syslog logging level defined for the profile. Assign a numeric identifier to log events based on criticality. Severity levels include 0 - <i>Emergency</i> , 1 - <i>Alert</i> , 2 - <i>Critical</i> , 3 - <i>Errors</i> , 4 - <i>Warning</i> , 5 - <i>Notice</i> , 6 - <i>Info</i> and 7 - <i>Debug</i> . The default logging level is 4.
Buffered Logging Level	Event severity coincides with the syslog logging level defined for the profile. Assign a numeric identifier to log events based on criticality. Severity levels include 0 - <i>Emergency</i> , 1 - <i>Alert</i> , 2 - <i>Critical</i> , 3 - <i>Errors</i> , 4 - <i>Warning</i> , 5 - <i>Notice</i> , 6 - <i>Info</i> and 7 - <i>Debug</i> . The default logging level is 4.
Time to Aggregate Repeated Messages	Define the increment (or interval) system events are logged on behalf of this controller profile. The shorter the interval, the sooner the event is logged. Either define an interval in <i>Seconds</i> (0 - 60) or <i>Minutes</i> (0 -1). The default value is 0 seconds.
Forward Logs to Controller	Select the checkbox to define a log level for forwarding event logs to the control. Log levels include <i>Emergency</i> , <i>Alert</i> , <i>Critical</i> , <i>Error</i> , <i>Warning</i> , <i>Notice</i> , <i>Info</i> and <i>Debug</i> . The default logging level is <i>Error</i> .

10. Refer to the **System Event Messages** section to define or override how controller system messages are logged and forwarded on behalf of the controller profile.

Event System Policy	Select an Event System Policy from the pull-down menu. If an appropriate policy does not exist click the create button to make a new policy.
Enable System Events	Select the Enable System Events check box to allow the controller profile to capture system events and append them to a log file. It's important to log individual events to discern an overall pattern that may be negatively impacting controller performance. This setting is enabled by default.
Enable System Event Forwarding	Select the Enable System Event Forwarding box to enable the forwarding of system events to another controller or cluster member. This setting is enabled by default.

11. Refer to the **Events E-mail Notification** section to define or override how system event notification emails are sent.

SMTP Server	Specify either the Hostname or IP Address of the outgoing SMTP server where notification emails will be sent from.
Port of SMTP	If a non-standard SMTP port is used on the outgoing SMTP server check this box and specify a port between 1 and 65,535 for the outgoing SMTP server to use.
Sender Email Address	Specify the email address that notification emails will be sent from. This will be the from address on notification emails.
Username for SMTP Server	Specify the username of the sender on the outgoing SMTP server. Many SMTP servers require users to authenticate with an username and password before sending email through the server.
Password for SMTP Server	Specify the password associated with the username of the sender on the outgoing SMTP server. Many SMTP servers require users to authenticate with an username and password before sending email through the server.

12. Select **OK** to save the changes made to the profile's Management Settings. Select **Reset** to revert to the last saved configuration.

13. Select **Firmware** from the Management menu.

The screenshot shows the 'Profile Management Firmware' configuration window. It contains the following settings:

- Auto Install via DHCP Option:**
 - Enable Configuration Update:
 - Enable Firmware Update:
- Legacy Device Firmware Management:**
 - Migration Firmware from AP71xx: 4.x path:
- Automatic Adopted AP Firmware Upgrade:**
 - Enable Controller Upgrade of AP Firmware:
 - AP71xx
 - AP650
 - AP621
 - AP6521
 - AP6511
 - AP6532
 - Number of Concurrent Upgrades: 10 (1 to 20 APs)
- Legacy Settings:**
 - Legacy AP650 Auto Downgrade:
 - Enable Update Legacy Device Firmware:

Buttons at the bottom right: OK, Reset, Exit.

FIGURE 275 Profile Management Firmware screen

14. Refer to the **Auto Install via DHCP Option** section to configure automatic configuration file and firmware updates.

Enable Configuration Update Select the Enable Configuration Update radio button (from within the Automatic Configuration Update field) to enable automatic configuration file updates for the controller profile from a location external to the controller. If enabled (the setting is disabled by default), provide a complete path to the target configuration file used in the update.

Enable Firmware Upgrade Select this option to enable automatic controller firmware upgrades (for this controller profile) from a user defined remote location. This value is disabled by default.

15. Refer to the **Legacy Device Firmware Management** field to define or whether Brocade Mobility 650 Access Point and AP-7131 model devices can upgrade to newer firmware versions or downgrade to legacy firmware versions.

Migration Firmware from AP7131 4.x path Provide a complete path to the target firmware used to support a legacy AP-7131 firmware update. The length of the path cannot exceed 253 characters.

16. Use the parameters within the **Automatic Adopted AP Firmware Upgrade** section to define an automatic firmware upgrade from a controller based file.

Enable Controller Upgrade of AP Firmware Select this radio button to enable adopted Access Point radios to upgrade to a newer firmware version using its associated controller's most recent resident firmware file for that AP model. This parameter is disabled by default.

Number of Concurrent Upgrades. Use the spinner control to define the maximum number (1 - 20) of adopted APs that can receive a firmware upgrade at the same time. Keep in mind that during a firmware upgrade, the AP is offline and unable to perform its normal wireless client support function until the upgrade process is complete.

17. Refer to the parameters within the **Legacy Settings** section to configure settings for legacy devices.

Legacy AP650 Auto Downgrade Select this box to enable automatic downgrading of legacy AP650 Access Point firmware.

Enable Update Legacy Device Firmware Select this box to update legacy device firmware when connected to the controller.

18. Select **OK** to save the changes made to the profile's Management Firmware configuration. Select **Reset** to revert to the last saved configuration.

19. Select the **Heartbeat** option from the Management menu.




FIGURE 276 Profile Management Heartbeat screen

20. Select the **Service Watchdog** option to implement heartbeat messages to ensure other associated devices are up and running and capable of effectively interoperating with the controller. The Service Watchdog is enabled by default.
21. Select **OK** to save the changes made to the profile maintenance Heartbeat tab. Select **Reset** to revert to the last saved configuration.

Profile Management Configuration and Deployment Considerations

Profile Management Configuration

Before defining a profile's management configuration, refer to the following deployment guidelines to ensure the profile configuration is optimally effective:

- Define controller profile management access configurations providing both encryption and authentication. Management services like HTTPS, SSH and SNMPv3 should be used when possible, as they provide data privacy and authentication.
- Brocade Mobility recommends SNMPv3 be used for management profile configurations, as it provides both encryption and authentication.

Advanced Profile Configuration

A profile's advanced configuration is comprised of defining its MINT protocol configuration and the profile's NAS identifier and port ID attributes. MINT provides secure controller profile communications at the transport layer. Using MINT, a device can be configured to only communicate with other authorized (MINT enabled) devices. Therefore, MINT is well designed for controller profile support, wherein a group of managed devices share the same configuration attributes.

Refer to the advanced profile's Miscellaneous menu item to set the profile's NAS configuration. The profile database on the RADIUS server consists of user profiles for each connected *network access server* (NAS) port.

To set a profile's advanced configuration:

1. Select the Configuration tab from the Web UI.
2. Select **Profiles** from the Configuration tab.
3. Select **Manage Profiles** from the Configuration > Profiles menu.
4. Select **Advanced** and expand the menu item.

The following sub menu items are available as advanced profile configuration options:

- [Configuring MINT](#)
- [Advanced Profile Miscellaneous Configuration](#)

Configuring MINT

MINT provides the means to secure controller profile communications at the transport layer. Using MINT, a device can be configured to only communicate with other authorized (MINT enabled) devices.

managed devices can communicate with each other exclusively over a MINT security domain. Keys can also be generated externally using any application (like openssl). These keys must be present on the managed device managing the domain for key signing to be integrated with the UI. A MAP device that needs to communicate with another first negotiates a security context with that device. The security context contains the transient keys used for encryption and authentication. A secure network requires users to know about certificates and PKI. However, administrators do not need to define security parameters for Access Points to be adopted (secure WISPe being an exception, but that isn't a commonly used feature). Also, users can replace any device on the network or move devices around and they continue to work. Default security parameters for MINT are such that these scenarios continue to function as expected, with minimal user intervention required only when a new network is deployed.

To define a controller profile's MINT configuration:

1. Select **MINT Protocol** from the Advanced Profile's menu item.

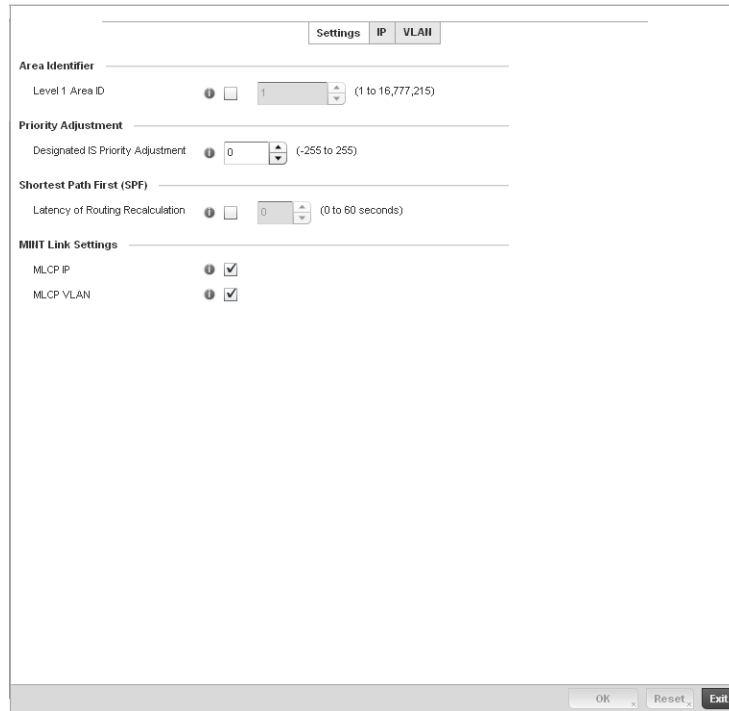


FIGURE 277 Advanced Profile MINT screen - Settings tab

2. The **Settings** tab displays by default.
3. Refer to the **Area Identifier** field to define the Level 1 and Level 2 Area IDs used by the profile's MINT configuration.

Level 1 Area ID Select the check box to enable a spinner control for setting the Level 1 Area ID between 1 - 4,294,967,295. The default value is disabled.

4. Define the following **Device Heartbeat Settings** in respect to devices supported by the controller profile:

Designated IS Priority Adjustment Use the spinner control to set a **Designated IS Priority Adjustment** setting between -255 and 255. This is the value added to the base level DIS priority to influence the *Designated IS* (DIS) election. A value of +1 or greater increases DISiness. The default setting is 0.

5. Select the **Latency of Routing Recalculation** check box (within the **Shortest Path First (SPF)** field) to enable the spinner control used for defining a latency period between 0 - 60 seconds. The default setting has the check box disabled.

The screenshot shows the 'Add IP MINT Link' configuration window. It contains the following fields and values:

- IP:** 192.68.38.15
- Port:** 1 (range: 1 to 65,535)
- Routing Level:** 1 (range: 1 to 2)
- Listening Link:** 1 (range: 0 to 1)
- Forced Link:**
- Link Cost:** 100 (range: 1 to 10,000)
- Hello Packet Interval:** 15 Seconds (range: 1 to 120)
- Adjacency Hold Time:** 46 Seconds (range: 2 to 600)

Buttons at the bottom: OK, Reset, Exit.

FIGURE 279 Advanced Profile MINT screen - IP Add tab

10. Set the following **Link IP** parameters to complete the MINT network address configuration:

IP	Define or override the IP address used by peer controllers for interoperation when supporting the MINT protocol.
Port	To specify a custom port for MiNT links, check this box and use the spinner control to define or override the port number between 1 and 65,535.
Routing Level	Use the spinner control to define or override a routing level of either 1 or 2.
Listening Link	Specify a listening link of either 0 or 1. UDP/IP links can be created by configuring a matching pair of links, one on each end point. However, that is error prone and doesn't scale. So UDP/IP links can also listen (in the TCP sense), and dynamically create connected UDP/IP links when contacted. The typical configuration is for the controller to have a listening UDP/IP link on the switch's IP address S.S.S.S, and for all the APs to have a regular UDP/IP link to S.S.S.S.
Forced Link	Check this box to specify the MiNT link as a forced link.
Link Cost	Use the spinner control to define or override a link cost between 1 - 10,000. The default value is 100.
Hello Packet Interval	Set or override an interval in either <i>Seconds</i> (1 - 120) or <i>Minutes</i> (1 - 2) for the transmission of hello packets. The default interval is 15 seconds.
Adjacency Hold Time	Set or override a hold time interval in either <i>Seconds</i> (2 - 600) or <i>Minutes</i> (1 - 10) for the transmission of hello packets. The default interval is 46 seconds.

11. Select the **VLAN** tab to display the link IP VLAN information shared by the devices managed by the controller's MINT configuration.

Link Cost	Use the spinner control to define or override a link cost between 1 - 10,000. The default value is 100.
Hello Packet Interval	Set or override an interval in either <i>Seconds</i> (1 - 120) or <i>Minutes</i> (1 - 2) for the transmission of hello packets. The default interval is 15 seconds.
Adjacency Hold Time	Set or override a hold time interval in either <i>Seconds</i> (2 - 600) or <i>Minutes</i> (1 - 10) for the transmission of hello packets. The default interval is 46 seconds.

14. Select **OK** to save the updates and overrides to the MINT Protocol configuration. Select **Reset** to revert to the last saved configuration.

Advanced Profile Miscellaneous Configuration

[Advanced Profile Configuration](#)

Refer to the advanced profile's Miscellaneous menu item to set the profile's NAS configuration. The profile database on the RADIUS server consists of user profiles for each connected *network access server* (NAS) port. Each profile is matched to a username representing a physical port. When the wireless controller authorizes users, it queries the user profile database using a username representative of the physical NAS port making the connection.

1. Select **Miscellaneous** from the Advanced Profile's menu item.

FIGURE 282 Advanced Profile Miscellaneous screen

2. Set a **NAS-Identifier Attribute** up to 253 characters in length.
This is the RADIUS NAS-Identifier attribute that typically identifies the Access Point or controller of controller where a RADIUS message originates.
3. Set a **NAS-Port-Id Attribute** up to 253 characters in length.

This is the RADIUS NAS port ID attribute which identifies the device port where a RADIUS message originates.

4. Select the **Capable** check box (within the RF Domain Manager field) to designate this specific profile managed device as being capable of being the RF Domain manager for a particular RF Domain. The default value is enabled.
5. Select the **Priority** check box (within the RF Domain Manager field) to set a priority value for this specific profile managed device. Once enabled, use the spinner control to set a device priority between 1 - 10,000. The higher the number set, the higher the priority in the RF Domain manager election process.
6. Select **OK** to save the changes made to the profile's Advanced Miscellaneous configuration. Select **Reset** to revert to the last saved configuration.

RF Domain Configuration

In this chapter

- [About RF Domains](#) 407
- [Managing RF Domains](#) 408

About RF Domains

A wireless controller configuration is composed of numerous elements including RF Domains, profiles, policies, WLANs and device specific configurations. RF Domains are used to assign regulatory, location and relevant policies to controllers. RF Domains are required, as each controller must be assigned at least one default RF Domain.

RF Domains allow administrators to assign configuration data to multiple devices deployed in a common coverage area, such as in a floor, building or site. Each RF Domain contains policies that can determine a Smart RF or WIPS configuration.

RF Domains enable administrators to override WLAN SSID name and VLAN assignments. This enables the deployment of a global WLAN across multiple sites and unique SSID name or VLAN assignments to groups of Access Points servicing the global WLAN. This new WLAN override technique eliminates the requirement for defining and managing a large number of individual WLANs and profiles.

A controller configuration contains (at a minimum) one default RF Domain and can optionally use additional user defined RF Domains:

- *Default RF Domain* - Automatically assigned to each controller and associated Access Point by default.
- *User Defined RF Domains* - Created by administrators and manually assigned to individual controllers or Access Points, but can be automatically assigned to Access Points using adoption policies.

Each controller and Access Point is assigned to only one RF Domain at a time. However, a user defined RF Domain can be assigned to multiple controllers or Access Points as required. User defined RF Domains can be manually assigned to controllers and Access Points or automatically assigned to Access Points using an AP provisioning policy.

Default RF Domains

Each controller utilizes a default RF Domain. Access Points are assigned to this default RF Domain as they are discovered by the controller. The default RF Domain can be used for single site deployments, where regional, regulatory and RF policies are common between devices. When regional, regulatory or RF policies need to be device specific, user defined RF Domains are recommended.

A default RF Domain can also omit configuration parameters to prohibit regulatory configuration from automatically being inherited by devices as they are discovered by the controller. This is desirable in multi-site deployments with devices spanning multiple countries. Omitting specific configuration parameters eliminates the risk of an incorrect country code from being automatically assigned to a device.

User Defined RF Domains

Configure and deploy user defined RF Domains for single or multiple sites when controllers and Access Points require unique regulatory and regional configurations, or unique Smart RF and WIPS policies. User defined RF Domains can be used to:

- Assign unique Smart RF or WIPS policies to Access Points deployed on different floors or buildings within in a site.
- Assign unique regional or regulatory configurations to controllers and Access Points deployed in different states or countries.
- Assign unique WLAN SSIDs and/or VLAN IDs to sites assigned a common WLAN without having to define individual WLANs for each site.

User defined RF Domains must be manually assigned to controllers, but can be manually or automatically assigned to Access Points. Manual RF Domain assignment can be performed using the controller CLI or UI applet by modifying each device's individual configuration and assigning a specific RF Domain to the device. Automatic RF Domain assignments can be made using an AP provisioning policy which can assign specific RF Domains to Access Points based on an Access Points model, serial number, VLAN, DHCP option, IP address or MAC address.

Automatic RF Domain assignments are useful in large deployments, as they enable plug-n-play Access Point deployments by automatically applying RF Domains to remote Access Points.

Managing RF Domains

Managing RF Domains entails configuring individual RF Domains as required and managing them as a collective set.

To review the configurations of existing RF Domains:

1. Select **Configuration > RF Domains** from the Web UI

The **RF Domain** screen displays within the main portion of the controller Web UI, and the **RF Domain Browser** displays in the lower, left-hand, portion of the controller Web UI.

2. Refer to the RF Domain screen to review high-level configuration data for existing RF Domain policies.

RF Domain	Location	Contact	Time Zone	Country
RF-Domain1	location1	contact1	EtcUTC	United States-us
RF-Domain2	location2	contact2		
RF-Domain3	location3	contact2		
rfdomain4				

Type to search in tables Row Count: 4

Add **Edit** **Delete**

FIGURE 283 RF Domain screen

3. Use the following (read only) information to determine whether a new RF Domain policy requires creation, or an existing RF Domain requires edit or deletion:

- RF Domain** Lists each policy's name, as assigned when it was created. The RF Domain name cannot be changed as part of the edit process. Only one RF Domain can be assigned to a controller or Access Point.
- Location** Displays the physical location assigned to the RF Domain policy. This name could be as specific as the floor of a building, or as generic as an entire site. The location defines the physical area where a common set of devices are deployed using the policy's RF Domain configuration.
- Contact** Lists the contact (or administrator) assigned to respond to events created by or impacting the RF Domain.
- Time Zone** Displays the geographic time zone set for each RF Domain policy. RF Domains can contain unique country codes and time zone information to controller and Access Points deployed across different states or countries, thus making them ideal for managing device configurations across different geographical deployments.
- Country Code** Display the two-digit country code set for the policy. The country code must be set accurately to avoid illegal operation, as device radios transmit in specific channels unique to their country of operation.

4. Refer to the **RF Domain Browser** to expand each existing RF Domain policy and review the device MAC addresses operating within the location defined and are using the configuration defined for the policy.

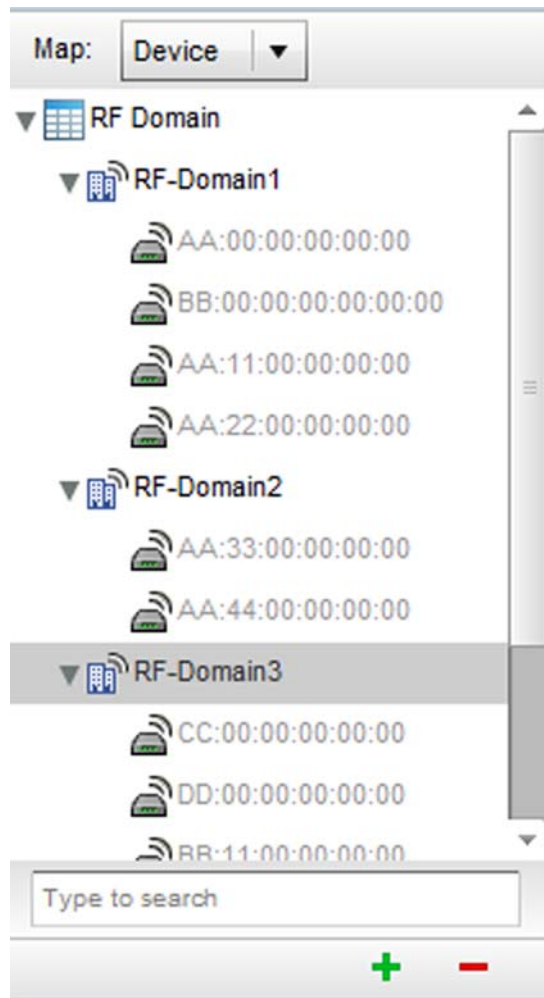


FIGURE 284 RF Domain Browser

5. Once the data within the RF Domain screen and RF Domain Browser is reviewed, determine whether a new policy requires creation, or if an existing policy requires edit or deletion. The management of RF Domains entails the following:
 - [RF Domain Basic Configuration](#)
 - [RF Domain Sensor Configuration](#)
 - [RF Domain Overrides](#)

RF Domain Basic Configuration

To set a RD Domain basic configuration:

1. Select **Configuration** > **RF Domains** from the Web UI
2. From the RF Domain screen, either select the **Add** button or highlight an existing RF Domain and select **Edit**. An RF Domain configuration can be permanently removed by highlighting it from the list and selecting **Delete**.

An existing RF Domain can also be modified by selecting it directly from the RF Domain Browser.

If adding or modifying an existing RF Domain, the RF Domain **Basic Configuration** screen displays by default.

Basic Configuration

Location: location1

Contact: contact1

Time Zone: (GMT) Etc/UTC

Country: United States-us

VLAN for Control Traffic: 1 (1 to 4,094)

SMART RF

SMART RF Policy: smartrf1

Enable Dynamic Channel:

2.4 GHz Channels: [Dropdown]

5 GHz Channels: [Dropdown]

2.4 GHz Radios: [Dropdown]

5 GHz Radios: [Dropdown]

Wireless IPS

WIPS Policy: <none>

Statistics

OK, Reset, Exit

FIGURE 285 RF Domain - Basic Configuration screen

3. Define the following **Basic Configuration** parameters for the RF Domain:

RF Domain

If creating a new RF Domain, assign it a name representative of its intended function. The name cannot exceed 32 characters. The name cannot be changed as part of the edit process.

Location

Assign the physical location of the controller RF Domain. This name could be as specific as the floor of a building, or as generic as an entire site. The location defines the physical area where a common set of device configurations are deployed and managed by the RF Domain policy.

Contact

Provide the name of the contact (or administrator) assigned to respond to events created by or impacting the RF Domain.

Time Zone	Displays the geographic time zone set for each RF Domain policy. RF Domains can contain unique country codes and time zone information to controller and Access Points deployed across different states or countries, thus making them ideal for managing device configurations across different geographical deployments.
Country	Define the two-digit country code set for the RF Domain. The country code must be set accurately to avoid the policy's illegal operation, as device radios transmit in specific channels unique to the country of operation.
VLAN for Traffic Control	Select the check box to enable a spinner control used for specifying the VLAN (within a range of 1 - 4,094) used for traffic control within this RF Domain.

When a radio fails or is faulty, a Smart RF policy can be used to provide automatic recovery by instructing neighboring Access Points to increase their transmit power to compensate for the coverage loss.

Once correct Access Point placement has been established, Smart-RF can optionally be leveraged for automatic detector radio selection. Smart-RF uses detector radios to monitor RF events and can be used to ensure adequate detector coverage is available. Manual detector radio selection can also be made using visualizations from the Brocade Mobility LANPlanner tool.

For an overview of Smart RF and instructions on how to create a Smart RF policy that can be used with a RF Domain, see *Smart RF Policy on page 6-301*.

4. Define the following **SMART RF** parameters for the RF Domain:

SMART RF Policy	Assign an existing Smart RF Policy to the RF Domain, or if none exist create a new one. Use the Smart RF Policy drop-down menu to navigate to existing Smart RF policies and select the one best suited to the function of the RF Domain. If none exist, select the Create icon and provide the required parameters to define a Smart RF configuration that can be used with the RF Domain. An existing policy can be edited by selecting the policy from the drop-down menu and selecting the Edit icon.
Enable Dynamic Channel	Check this box to enable dynamic channel switching for Smart RF radios.
2.4GHz Channels	Select channels from the pull-down menu and click the down arrow to move it to the list of channels used for 2.4GHz Smart RF radios.
5GHz Channels	Select channels from the pull-down menu and click the down arrow to move it to the list of channels used for 5GHz Smart RF radios.

5. Assign an existing **Wireless IPS (WIPS)** policy to the RF Domain, or if none exist create a new one.

Use the **WIPS Policy** drop-down menu to navigate to existing WIPS policies and select the one best suited to the function of the RF Domain. If none exist, select the **Create** icon and provide the required parameters to define a WIPS configuration that can be used with the RF Domain. An existing policy can be edited by selecting the policy from the drop-down menu and selecting the **Edit** icon.

A WIPS policy provides protection against wireless threats and acts as a key layer of security complementing wireless VPNs, encryption and authentication. A WIPS policy uses a dedicated sensor for actively detecting and locating rogue AP devices. After detection, WIPS uses mitigation techniques to block the devices by manual termination, air lockdown, or port suppression.

For an overview of WIPS and instructions on how to create a WIPS policy that can be used with a RF Domain, see [Configuring a WIPS Policy on page 9-447](#).

6. Refer to the **Statistics** field to define how RF Domain stats and updated.

NoC Update Interval	Set a NoC Update interval of 0, or between 5-3600 seconds for updates from the RF Domain manager to the controller.
Update Interval	Set a statistics Update interval of 0 or between 5-3600 seconds for updates retrieved from the wireless controller.

7. Refer to the Statistics field to set the following:

Window Index	Use the spinner control to set a numerical index used as an identifier for each RF Domain statistics defined.
Sample Interval	Use the spinner control to define the interval (in seconds) used by the controller to capture windowed statistics supporting the listed RF Domain configuration. The default is 5 seconds.
Window Size	Use the spinner control to set the number of samples used by the controller to define RF Domain statistics. The default value is 6 samples.

8. Select **OK** to save the changes to the Basic Configuration, or select **Reset** to Revert to the last saved configuration.

RF Domain Sensor Configuration

The Brocade Mobility *Wireless Intrusion Protection System* (WIPS) protects the managed network, wireless clients and Access Point radio traffic from attacks and unauthorized access. WIPS provides tools for standards compliance and around-the-clock wireless network security in a distributed environment. WIPS allows administrators to identify and accurately locate attacks, rogue devices and network vulnerabilities in real time and permits both a wired and wireless lockdown of wireless device connections upon acknowledgement of a threat.

In addition to dedicated Brocade Mobility AirDefense sensors, an Access Point radio can function as a sensor and upload information to a dedicated WIPS server (external to the controller). Unique WIPS server configurations can be used by RF Domains to ensure a WIPS server configuration is available to support the unique data protection needs of individual RF Domains.

WIPS is not supported on a WLAN basis, rather sensor functionality is supported on the Access Point radio(s) available to each managed WLAN. When an Access Point radio is functioning as a WIPS sensor, it's able to scan in sensor mode across all legal channels within 2.4 and 5.0 GHz. Sensor support requires a AirDefense WIPS Server on the network. Sensor functionality is not provided by the Access Point alone. The Access Point works in conjunction with a dedicated WIPS server.

To define a WIPS server configuration used with a RF Domain:

1. From the RF Domain screen, either select the **Add** button or highlight an existing policy and select **Edit**.
An existing policy can also be modified by selecting it directly from the RF Domain Browser.
2. Select the **Sensor Configuration** item from within the RF Domain screen.

Sensor Appliance Configuration

Server Id	IP Address	Port	
1	192.158.91.58	443	

Add Row

OK Reset Exit

FIGURE 286 RF Domain - Sensor WIPS screen

3. Either select the **+ Add Row** button to create a new WIPS server configuration or highlight an existing Sensor Server Configuration and select the **Delete** icon to remove it.
4. Use the spinner control to assign a numerical **Server ID** to each WIPS server defined. The server with the lowest defined ID is the first reached by the controller. The default ID is 1.
5. Provide the numerical (non DNS) **IP Address** of each server used as a WIPS sensor server by the RF Domain.
6. Use the spinner control to specify the **Port** of each WIPS server. The default port is 443.
7. Select **OK** to save the changes to the AirDefense WIPS configuration, or select **Reset** to Revert to the last saved configuration.

RF Domain Overrides

Within the managed network, each WLAN provides associated wireless clients with a *Service Set Identifier* (SSID). This has limitations because it requires wireless clients associate with different SSIDs to obtain QoS and security policies. However, a Brocade Mobility managed RF Domain can have WLANs assigned and advertise a single SSID, but allow users to inherit different QoS or security policies. Use the Override SSID screen to assign WLANs an override SSID as needed for the RF Domain.

The controller allows the mapping of a WLAN to more than one VLAN. When a wireless client associates with a WLAN, it is assigned a VLAN in such a way that users are load balanced across VLANs. The VLAN is assigned from the pool representative of the WLAN. The controller tracks the number of client users per VLAN, and assigns the least used/loaded VLAN to the wireless client. This number is tracked on a per-WLAN basis.

To define an override SSID and override VLAN configuration to used with a RF Domain:

1. From the RF Domain screen, either select the **Add** button or highlight an existing policy and select **Edit**.

An existing policy can also be modified by selecting it directly from the RF Domain Browser.

2. Select the **Overrides** item from within the RF Domain screen.

WLAN	SSID	
vLAN102	eng1	🗑️

Add Row

OK Reset Exit

FIGURE 287 RF Domain Override SSID screen

The Overrides screen is partitioned into two tabs, with the **Override SSID** screen displayed by default.

3. Either select the **+ Add** button to create a new Override SSID configuration. Highlight an existing Sensor Server Configuration and select the Delete icon to remove it from the table.
4. Use the **WLAN** drop-down menu to select a existing WLAN to be supplied an override SSID.

If a WLAN configuration has not been defined, you'll need to select the **Create** button and define at least one complete WLAN configuration. For detailed information on the steps required to create a WLAN, see *Wireless LAN Policy on page 6-228*.
5. Enter the name of the **SSID** to use as the override SSID.
6. Select **OK** to save the changes to the Override SSID configuration, or select **Reset** to Revert to the last saved configuration.
7. Select the **Override VLAN** tab.

The Override VLAN screen lists those WLANs available for override.

WLAN	VLAN
WLAN01	1

Type to search in tables

Row Count: 1

Add Edit Delete Exit

FIGURE 288 RF Domain Override VLAN screen

8. Either select **Add** to define a new VLAN override configuration, choose an existing WLAN and select **Edit** to change the override VLAN and limit or select **Delete** to remove a WLAN's override VLAN configuration.

VLAN	Wireless Client Limit	
1	64	

Add Row

OK Reset Exit

FIGURE 289 RF Domain Override VLAN Add screen

9. Use the **VLAN** spinner control to change the add additional VLANs for the selected WLAN. By default, VLAN 1 is configured for any selected WLAN.
10. Use the **Wireless Client Limit** spinner control to set the client user limit for the VLAN. The maximum allowed client limit is 8192 per VLAN. VLANs can be defined between 1 - 4094. The default setting is 0.

11. Select **OK** to save the changes to the Override VLAN configuration, or select **Reset** to Revert to the last saved configuration.

RF Domain Deployment Considerations

Before defining RF Domain policies, refer to the following deployment guidelines to ensure the configurations are optimally effective:

- Each controller utilizes a default RF Domain. Access Points are assigned to this default RF Domain as they are discovered by the controller. The default RF Domain can be used for single site deployments, where regional, regulatory and RF policies are common between devices.
- User defined RF Domains must be manually assigned to controllers, but can be manually or automatically assigned to Access Points.
- A Rogue AP detection configuration is a central component of an RF Domain policy, as it provides the RF Domain policy with the means to filter potentially threatening devices from operating with devices approved within the managed network.
- WIPS is not supported on a WLAN basis, rather sensor functionality is supported on the Access Point radio(s) available to each managed WLAN.
- When planning sensor coverage, a minimum of 1 detector radio is recommended per 4 Access Points deployed. To ensure effective placement, Brocade Mobility' LANPlanner can be used to provide predictive planning services and visualization to ensure adequate radio coverage is provided based on site application and device requirements. LANPlanner provides visualization tools ensuring adequate radio coverage for client radios and sensors. A physical site survey should also be performed to verify client radio coverage, before a final deployment.
- Both default and user defined RF Domains contain policies and configuration parameters. Changes made to policies or configuration parameters are automatically inherited by all the controllers and Access Points assigned to the RF Domain.

Security Configuration

In this chapter

- [Wireless Firewall](#) 419
- [Wireless Client Roles](#) 438
- [Intrusion Prevention](#) 446

When taking precautions to secure wireless traffic from a client to an access points and wireless controller, the network administrator should not lose sight of the security solution in its entirety, since the chain is as weak as its weakest link. A wireless controller managed network provides seamless data protection and user validation to protect and secure data at each vulnerable point in the network. The controller supports a Layer 2 wired/wireless Firewall and *Wireless Intrusion Protection System* (WIPS) capabilities at the WLAN, while additionally strengthened with a premium multi-vendor overlay Security Solution from Air Defense with 24x7 dedicated protection. This security is offered at the most granular level, with role-based and location based secure network access control available to users based on identity as well as the security posture of the client device. For more information, see:

- [Wireless Firewall](#)
- [Wireless Client Roles](#)
- [Intrusion Prevention](#)

Wireless Firewall

A Firewall is a mechanism enforcing access control, and is considered a first line of defense in protecting proprietary information within the wireless controller managed network. The means by which this is accomplished varies, but in principle, a Firewall can be thought of as mechanisms both blocking and permitting data traffic within the wireless controller managed network. Firewalls implement uniquely defined access control policies, so if you don't have an idea of what kind of access to allow or deny, a Firewall is of little value, and in fact could provide a false sense of network security.

With Brocade Mobility wireless controllers, Firewalls are configured to protect against unauthenticated logins from outside the wireless controller managed network. This helps prevent hackers from accessing wireless clients within the wireless controller managed network. Well designed Firewalls block traffic from outside the wireless controller managed network, but permit wireless controller authorized users to communicate freely with outside the wireless controller managed network.

Firewalls can be implemented in both hardware and software, or a combination of both. All messages entering or leaving the wireless controller pass through the Firewall, which examines each message and blocks those not meeting the security criteria (rules) defined.

Firewall rules define the traffic permitted or denied within the wireless controller managed network. Rules are processed by a Firewall device from first to last. When a rule matches the network traffic a wireless controller is processing, the Firewall uses that rule's action to determine whether traffic is allowed or denied.

Rules comprise conditions and actions. A condition describes a traffic stream of packets. Define constraints on the source and destination device, the service (for example, protocols and ports), and the incoming interface. An action describes what should occur to packets matching the conditions set. For example, if the packet stream meets all conditions, traffic is permitted, authenticated and sent to the destination device.

Additionally, IP and MAC rule based Firewall filtering can be deployed to apply Firewall policies to traffic being bridged by centrally managed radios. IP and MAC filtering can be employed to permit or restrict traffic exchanged between hosts, hosts residing on separate WLANs or hosts forwarding traffic to wired devices.

For more information, refer to the following:

- [Configuring a Firewall Policy](#)
- [Configuring IP Firewall Rules](#)
- [Configuring MAC Firewall Rules](#)
- [Firewall Deployment Considerations](#)

Configuring a Firewall Policy

[Wireless Firewall](#)

To configure a Firewall on the wireless controller:

Select **Configuration > Security > Wireless Firewall > Firewall Policy** to display existing Firewall policies.

The **Wireless Firewall** screen lists existing Firewall policies. An existing policy can be selected and applied to the controller. The user has the option of displaying the configurations of each Wireless Firewall Policy, or referring to the **Wireless Firewall Browser** and selecting individual policies for review.

2. Select **Add** to create a new Wireless Firewall policy. Select an existing policy and click **Edit** to modify the attributes of that policy.
3. The **Denial of Services** tab displays by default.
4. When adding a new policy, first enter a name in the Firewall Policy box. The name must not exceed 64 characters. Once a name has been specified, click **OK** to enable the other parameters within the screen.

The Wireless Firewall Policy configuration is divided into the following tabs:

- [Firewall Policy Denial of Service](#)
- [Firewall Policy Storm Control](#)
- [Firewall Policy Advanced Settings](#)

Firewall Policy Denial of Service

[Adding and Editing Wireless Firewall Policies](#)

A *denial of service* (DoS) attack is an attempt to make a computer or network resource unavailable to its intended users. Although the means to carry out a DoS attack will vary, it generally consists of a concerted effort of one or more persons attempting to prevent a device, site or service from functioning temporarily or indefinitely.

Most DoS attacks involve saturating the target device with external communications requests so it cannot respond to legitimate traffic or respond so slowly the device becomes unavailable in respect to its defined data rate. DoS attacks are implemented by either forcing targeted devices to reset or consuming the devices resources so it can no longer provide service.

To define a Denial of Service configuration for a Firewall policy:

1. From the **Firewall Policy** configuration page, select the **Denial of Service** tab. The Denial of Service tab displays by default.

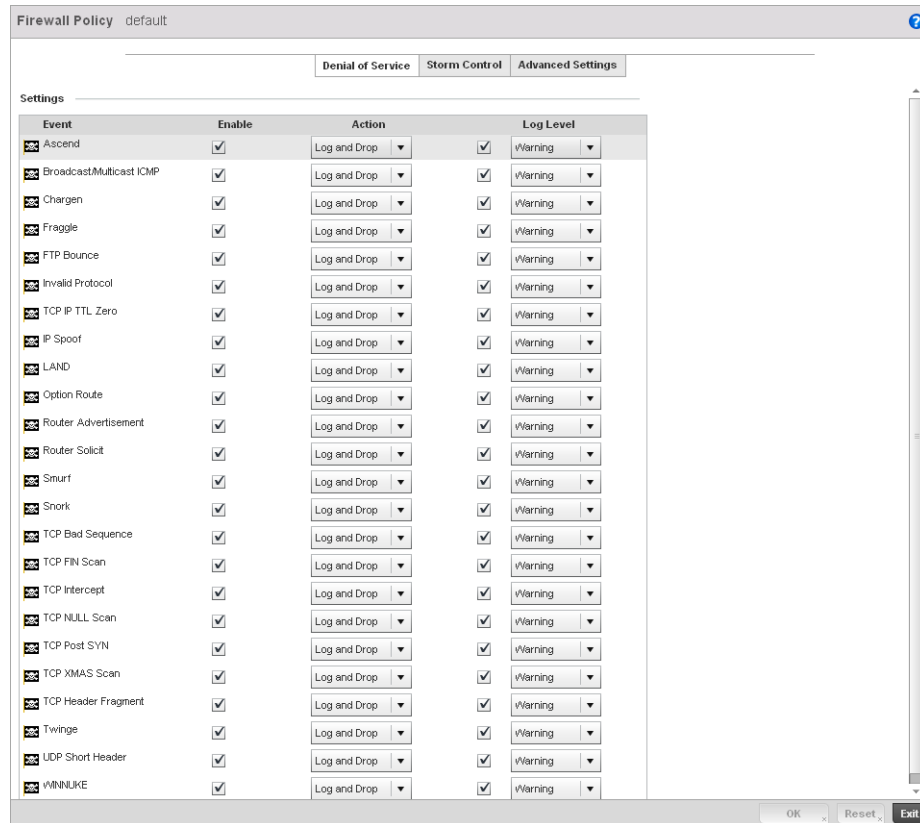


FIGURE 291 Wireless Firewall Add/Edit Denial of Service screen

2. The **Settings** window contains a list of all of the *Denial of Service* (DoS) attacks that the wireless controller's Firewall has filters for. Each DoS filter contains the following four items:

Event	The Event column lists the name of each Denial of Service attack.
Enable	Checking the Enable box will set the Firewall Policy to filter the associated Denial of Service attack based on the selection in the Action column.
Action	If a Denial of Service filter is enabled, chose an action from the drop-down menu to determine how the Firewall Policy treats the associated DoS attack. <i>Log and Drop</i> - An entry for the associated DoS attack is added to the log and then the packets are dropped. <i>Log Only</i> - An entry for the associated DoS attack is added to the log. No further action is taken. <i>Drop Only</i> - The DoS packets is dropped. No further action is taken.
Log Level	To enable logging to the system log, check the box in the Log Level column. Then select a standard Syslog level from the Log Level drop-down menu.

Denial of Service Attacks Table

Refer to the following for a summary of each Denial of Service attack the Firewall can filter.

Ascend	The Ascend DoS attacks are a series of attacks that target known vulnerabilities in various versions of Ascend routers.
Broadcast/Multicast ICMP	Broadcast or Multicast ICMP DoS attacks are a series of attacks that take advantage of ICMP behavior in response to echo replies. These usually involve spoofing the source address of the target and sending ICMP broadcast or multicast echo requests to the rest of the network and in the process flooding the target machine with replies.
Chargen	The Chargen attack establishes a Telnet connection to port 19 and attempts to use the character generator service to create a string of characters which is then directed to the DNS service on port 53 to disrupt DNS services.
Fraggle	The Fraggle DoS attack uses a list of broadcast addresses to send spoofed UDP packets to each broadcast address' echo port (port 7). Each of those addresses that have port 7 open will respond to the request generating a lot of traffic on the network. For those that do not have port 7 open they will send an unreachable message back to the originator, further clogging the network with more traffic.
FTP Bounce	The FTP Bounce DoS attack uses a vulnerability in the FTP "PORT" command as a way to scan ports on a target machine by using another machine in the middle.
Invalid Protocol	Attackers may use vulnerability in the endpoint implementation by sending invalid protocol fields, or may misuse the misinterpretation of endpoint software. This can lead to inadvertent leakage of sensitive network topology information, call hijacking, or a DoS attack.
TCP IP TTL Zero	The TCP IP TTL Zero DoS attack sends spoofed multicast packets onto the network which have a <i>Time To Live</i> (TTL) of 0. This causes packets to loop back to the spoofed originating machine, and can cause the network to overload.
IP Spoof	IP Spoof is a category of Denial of Service attack that sends IP packets with forged source addresses. This can hide the identity of the attacker.
LAND	The LAND DoS attack sends spoofed packets containing the SYN flag to the target destination using the target port and IP address as both the source and destination. This will either crash the target system or result in high resource utilization slowing down all other processes.
Option Route	Enables the IP Option Route denial of service check in the controller's firewall.
Router Advertisement	In this attack, the attacker uses ICMP to redirect the network router function to some other host. If that host can not provide router services, a DoS of network communications occurs as routing stops. This can also be modified to single out a specific system, so that only that system is subject to attack (because only that system sees the 'false' router). By providing router services from a compromised host, the attacker can also place themselves in a "man-in-the-middle" situation and take control of any open channel at will (as mentioned earlier, this is often used with TCP packet forgery and spoofing to intercept and change open TELNET sessions).

Router Solicit	<p>The ICMP Router Solicitation scan is used to actively find routers on a network. Of course, a hacker could set up a protocol analyzer to detect routers as they broadcast routing information on the network. In some instances, however, routers may not send updates. For example, if the local network does not have other routers, the router may be configured to not send routing information packets onto the local network.</p> <p>ICMP offers a method for router discovery. Clients send ICMP router solicitation multicasts onto the network, and routers must respond (as defined in RFC 1122). (For more information about the process of ICMP router solicitation, see "Routing Sequences for ICMP.")</p> <p>By sending ICMP Router Solicitation packets (ICMP type 9) on the network and listening for ICMP Router Discovery replies (ICMP type 10), hackers can build a list of all of the routers that exist on a network segment. Hackers often use this scan to locate routers that do not reply to ICMP echo requests</p>
Smurf	<p>The Smurf DoS Attack sends ICMP echo requests to a list of broadcast addresses in a row, and then repeats the requests, thus flooding the network.</p>
Snork	<p>The Snork DoS attack uses UDP packet broadcasts to consume network and system resources.</p>
TCP Bad Sequence	<p>Enables a TCP Bad Sequence denial of service check in the controller's firewall.</p>
TCP FIN Scan	<p>Hackers use the TCP FIN scan to identify listening TCP port numbers based on how the target device reacts to a transaction close request for a TCP port (even though no connection may exist before these close requests are made). This type of scan can get through basic firewalls and boundary routers that filter on incoming TCP packets with the Finish (FIN) and ACK flag combination. The TCP packets used in this scan include only the TCP FIN flag setting.</p> <p>If the target device's TCP port is closed, the target device sends a TCP RST packet in reply. If the target device's TCP port is open, the target device discards the FIN and sends no reply.</p>

TCP Intercept	<p>A SYN-flooding attack occurs when a hacker floods a server with a barrage of requests for connection.</p> <p>Because these messages have unreachable return addresses, the connections cannot be established. The resulting volume of unresolved open connections eventually overwhelms the server and can cause it to deny service to valid requests, thereby preventing legitimate users from connecting to a Web site, accessing email, using FTP service, and so on.</p> <p>The TCP intercept feature helps prevent SYN-flooding attacks by intercepting and validating TCP connection requests. In intercept mode, the TCP intercept software intercepts TCP synchronization (SYN) packets from clients to servers that match an extended access list. The software establishes a connection with the client on behalf of the destination server, and if successful, establishes the connection with the server on behalf of the client and knits the two half-connections together transparently. Thus, connection attempts from unreachable hosts will never reach the server. The software continues to intercept and forward packets throughout the duration of the connection. The number of SYNs per second and the number of concurrent connections proxied depends on the platform, memory, processor, and other factors. In the case of illegitimate requests, the software's aggressive timeouts on half-open connections and its thresholds on TCP connection requests protect destination servers while still allowing valid requests.</p> <p>When establishing a security policy using TCP intercept, you can choose to intercept all requests or only those coming from specific networks or destined for specific servers. You can also configure the connection rate and threshold of outstanding connections. Optionally operate TCP intercept in watch mode, as opposed to intercept mode. In watch mode, the software passively watches the connection requests flowing through the router. If a connection fails to get established in a configurable interval, the software intervenes and terminates the connection attempt.</p>
TCP Null Scan	<p>Hackers use the TCP NULL scan to identify listening TCP ports. This scan also uses a series of strangely configured TCP packets, which contain a sequence number of 0 and no flags. Again, this type of scan can get through some firewalls and boundary routers that filter incoming TCP packets with standard flag settings.</p> <p>If the target device's TCP port is closed, the target device sends a TCP RST packet in reply. If the target device's TCP port is open, the target discards the TCP NULL scan, sending no reply.</p>
TCP Post SYN	<p>A remote attacker may be attempting to avoid detection by sending a SYN frame with a different sequence number than the original SYN. This can cause an <i>Intrusion Detection System (IDS)</i> to become unsynchronized with the data in a connection. Subsequent frames sent during the connection are ignored by the IDS.</p>
TCP XMAS Scan	<p>The TCP XMAS Scan floods the target system with TCP packets including the FIN, URG, and PUSH flags. This is used to determine details about the target system and can crash a system.</p>
TCP Header Fragment	<p>Enables the TCP Header Fragment denial of service check in the controller's firewall.</p>
Twinge	<p>The Twinge DoS attack sends ICMP packets and cycles through using all ICMP types and codes. This can crash some Windows systems.</p>
UDP Short Header	<p>Enables the UDP Short Header denial of service check in the controller's firewall.</p>
WINNUKE	<p>The WINNUKE DoS attack sends a large amount of data to UDP port 137 to crash the NETBIOS service on windows and can also result on high CPU utilization on the target machine.</p>

3. Select **OK** to update the Denial of Service settings. Select **Reset** to revert to the last saved configuration.

Firewall Policy Storm Control

[Adding and Editing Wireless Firewall Policies](#)

The Firewall maintains a facility to control packet storms. Storms are packet bombardments that exceed the high threshold value configured for an interface. During a storm, packets are throttled until the rate falls below the configured rate, severely impacting performance for the RF Domain manager interface. Thresholds are configured in terms of packets per second.

To define a Storm Control configuration for a Firewall policy:

1. From the **Firewall Policy** configuration page, select the **Storm Control** tab.

The screenshot shows the 'Firewall Policy default' configuration window with the 'Storm Control' tab selected. At the top, there are three tabs: 'Denial of Service', 'Storm Control', and 'Advanced Settings'. Below the tabs, there are two main sections:

Storm Control Settings

Traffic Type	Interface Type	Interface Name	Packets per Second	

Storm Control Logging

Traffic Type	Logging	

At the bottom of the window, there are three buttons: 'OK', 'Reset', and 'Exit'.

FIGURE 292 Wireless Firewall Add/Edit Storm Control screen

2. Refer to the **Storm Control Settings** field to set the following:

- Traffic Type** Use the drop-down menu to define the traffic type for which the Storm Control configuration applies. Options include *ARP*, *Broadcast*, *Multicast* and *Unicast*.
- Interface Type** Use the drop-down menu to define the controller interface for which the Storm Control configuration is applied. Only the specified interface uses the defined filtering criteria. Options include *Ethernet*, *WLAN* and *Port Channel*.
- Interface Name** Use the drop-down menu to refine the interface selection to a specific WLAN or physical controller port. This helps with threshold configuration for potentially impacted controller interfaces.
- Packets per Second** Select the check box to activate the spinner control used for specifying the packets per second threshold for activating the Storm Control mechanism.

3. Select **+ Add Row** as needed to add additional Storm Control configurations for other traffic types or interfaces. Select the **Delete** icon as required to remove selected rows.
4. Refer to the **Storm Control Logging** field to define how storm events are logged for the controller.

Traffic Type

Use the drop-down menu to define the traffic type for which the Storm Control logging configuration applies. Options include *ARP*, *Broadcast*, *Multicast* and *Unicast*.

Logging

Select the check box to activate the spinner control used for specifying the standard log level used if a Storm Control attack is detected. The default log level is Warning.

5. Select **+ Add Row** as needed to add additional Storm Control log entries for other interfaces. Select the **Delete** icon as required to remove selected rows.
6. Select **OK** to update the Storm Control settings. Select **Reset** to revert to the last saved configuration.

Firewall Policy Advanced Settings

[Adding and Editing Wireless Firewall Policies](#)

To define a Firewall policy Advanced Configuration:

1. Select the **Advanced Settings** tab from the **Firewall Policy** configuration page.

The screenshot displays the 'Firewall Policy' configuration page, specifically the 'Advanced Settings' tab. The interface is organized into several sections:

- Enable Firewall:** Radio buttons for 'Enabled' (selected) and 'Disabled'.
- General:**
 - Enable Proxy ARP:
 - DHCP Broadcast to Unicast:
 - L2 Stateful Packet Inspection:
 - IPMAC Conflict Enable:
 - IPMAC Conflict Logging: Warning
 - IPMAC Conflict Action: Log and Drop
 - IPMAC Routing Conflict Enable:
 - IPMAC Routing Conflict Logging: Warning
 - IPMAC Routing Conflict Action: Log and Drop
 - DNS Snoop Entry Timeout: 1800 (30 to 86,400 second)
 - IP TCP Adjust MSS: 472 (472 to 1,460 bytes)
 - TCP MSS Clamping:
 - Max: Fragments/Datagram: 140 (2 to 8,129)
 - Max: Defragmentations/Host: 8 (1 to 16,384)
 - Min Length Required: 8 (8 to 1,500 bytes)
 - IPv4 Virtual Defragmentation:
- Application Layer Gateway:**
 - FTP ALG:
 - TFTP ALG:
 - SIP ALG:
 - DNS ALG:
- Firewall Enhanced Logging:**
 - Log Dropped ICMP Packets: None
 - Log Dropped Malformed Packets: None
 - Enable Verbose Logging:
- Stateful Flow Checks:**
 - Enable Stateful DHCP Checks:
- Flow Timeout:**
 - TCP Close Wait: 30 Seconds (1 to 32,400)
 - TCP Established: 3 Hours (0 to 9)
 - TCP Reset: 10 Seconds (1 to 32,400)
 - TCP Setup: 10 Seconds (1 to 32,400)
 - Stateless TCP Flow: 90 Seconds (1 to 32,400)
 - Stateless FIN/RESET Flow: 10 Seconds (1 to 32,400)
 - ICMP: 30 Seconds (1 to 32,400)
 - UDP: 90 Seconds (15 to 32,400)
 - Any Other Flow: 5 Minutes (0 to 540)
- TCP Protocol Checks:**
 - Check TCP states where a SYN packet tears down the flow
 - Check unnecessary resends of TCP packet
 - Check Sequence Number in ICMP Unreachable error packets
 - Check Acknowledgement Number in RST packets
 - Check Sequence Number in RST packets

At the bottom of the screen, there are buttons for 'OK', 'Reset', and 'Exit'.

FIGURE 293 Wireless Firewall Add/Edit Advanced Settings screen

2. Refer to the **Enable Firewall** radio buttons to define the Firewall as either *Enabled* or *Disabled*. The Firewall is enabled by default.

If disabling the Firewall, a confirmation prompt displays stating NAT, wireless hotspot, proxy ARP, deny-static-wireless-client and deny-wireless-client sending not permitted traffic excessively will be disabled.

3. Select **OK** to continue disabling the hotspot.
4. Refer to the **General** field to enable or disable the following Firewall configuration parameters:

Enable Proxy ARP	Select this check box to allow the Firewall Policy to use Proxy ARP responses for this policy on behalf of another device. Proxy ARP allows the Firewall to handle ARP routing requests for devices behind the Firewall. This feature is enabled by default.
DHCP Broadcast to Unicast	Select this check box to enable the conversion of broadcast DHCP offers to unicast. Converting DHCP broadcast traffic to unicast traffic can help reduce network traffic loads. This feature is disabled by default.
L2 Stateful Packet Inspection	Select the check box to enable stateful packet inspection for RF Domain manager routed interfaces within the Layer 2 Firewall. This feature is disabled by default.
IPMAC Conflict Enable	When multiple devices on the network have the same IP or MAC address this can create routing issues for traffic being passed through the Firewall. To avoid these issues, enable Conflict Detection to enable IP and MAC conflict detection. This feature is disabled by default.
IPMAC Conflict Logging	Select this option to enable logging for IP and MAC address conflict detection. This feature is disabled by default.
IPMAC Conflict Action	Use the drop-down menu to set the action taken when an attack is detected. Options include Log Only, Drop Only or Log and Drop. The default setting is Log and Drop.
IPMAC Routing Conflict Enable	Select this option to enable IPMAC Routing Conflict detection. This is also known as a Hole-196 attack in the network. This feature helps to detect if the client is sending routed packets to the correct router-mac-address.
IPMAC Routing Conflict Logging	Select enable logging for IPMAC Routing Conflict detection. This feature is disabled by default.
IPMAC Routing Conflict Action	Use the drop-down menu to set the action taken when an attack is detected. Options include Log Only, Drop Only or Log and Drop. The default setting is Log and Drop.
DNS Snoop Entry Timeout	Select this option and set a timeout, in seconds, for DNS Snoop Entry. DNS Snoop Entry stores information such as Client to IP Address and Client to Default Gateway(s) and uses this information to detect if the client is sending routed packets to a wrong MAC address.
IP TCP Adjust MSS	Select this option and adjust the value for the maximum segment size (MSS) for TCP segments on the router. Set a value between 472 bytes and 1,460 bytes to adjust the MSS segment size. The default value is 472 bytes.
TCP Adjust MSS	Select this option to enable TCP MSS Clamping. TCP MSS Clamping allows configuration for the maximum segment size of packets at a global level.
Max Fragments/Datagram	Set a value for the maximum number of fragments (between 2 and 8,129) allowed in a datagram before it is dropped. The default value is 140 fragments.

Max Defragmentations/Host	Set a value for the maximum number of defragmentations, between 1 and 16,384 allowed per host before it is dropped. The default value is 8.
Min Length Required	Select this option and set a minimum length, between 8 bytes and 1,500 bytes, to enforce a minimum packet size before being subject to fragment based attack prevention.
IPv4 Virtual Defragmentation	Select this option to enable IPv4 Virtual Defragmentation, this helps prevent IPv4 fragments based attacks such as tiny fragments or large number of ipv4 fragments.

5. The Firewall policy allows traffic filtering at the application layer using the **Application Layer Gateway** feature. The Application Layer Gateway provides filters for the following common protocols:

FTP ALG	Check this check box to allow FTP traffic through the Firewall using its default ports. This feature is enabled by default.
TFTP ALG	Check this check box to allow TFTP traffic through the Firewall using its default ports. This feature is enabled by default.
SIP ALG	Check this check box to allow SIP traffic through the Firewall using its default ports. This feature is enabled by default.
DNS ALG	Check the Enable box to allow DNS traffic through the Firewall using its default ports. This feature is enabled by default.

6. Refer to the **Firewall Enhanced Logging** field to set the following parameters:

Log Dropped ICMP Packets	Use the drop-down menu to define how dropped ICMP packets are logged. Logging can be rate limited for one log instance every 20 seconds. Options include <i>Rate Limited</i> , <i>All</i> or <i>None</i> . The default setting is <i>None</i> .
Log Dropped Malformed Packets	Use the drop-down menu to define how dropped malformed packets are logged. Logging can be rate limited for one log instance every 20 seconds. Options include <i>Rate Limited</i> , <i>All</i> or <i>None</i> . The default setting is <i>None</i> .
Enable Verbose Logging	Check this box to enable verbose logging mode for the firewall.

7. Select the **Enable Stateful DHCP Checks** check box to enable the stateful checks of DHCP packet traffic through the Firewall. The default setting is enabled. When enabled, all DHCP traffic flows are inspected.

8. Define **Flow Timeout** intervals for the following flow types impacting the Firewall:

TCP Close Wait	Define a flow timeout value in either <i>Seconds</i> (1 - 32,400), <i>Minutes</i> (1 - 540) or <i>Hours</i> (1 - 9). The default setting is 30 seconds.
TCP Established	Define a flow timeout value in either <i>Seconds</i> (1 - 32,400), <i>Minutes</i> (1 - 540) or <i>Hours</i> (1 - 9). The default setting is 10,800 seconds.
TCP Reset	Define a flow timeout value in either <i>Seconds</i> (1 - 32,400), <i>Minutes</i> (1 - 540) or <i>Hours</i> (1 - 9). The default setting is 10 seconds.
TCP Setup	Define a flow timeout value in either <i>Seconds</i> (1 - 32,400), <i>Minutes</i> (1 - 540) or <i>Hours</i> (1 - 9). The default setting is 10 seconds.
Stateless TCP Flow	Define a flow timeout value in either <i>Seconds</i> (1 - 32,400), <i>Minutes</i> (1 - 540) or <i>Hours</i> (1 - 9). The default setting is 90 seconds.

Stateless FIN/RESET Flow	Define a flow timeout value in either <i>Seconds</i> (1 - 32,400), <i>Minutes</i> (1 - 540) or <i>Hours</i> (1 - 9). The default setting is 10 seconds.
ICMP	Define a flow timeout value in either <i>Seconds</i> (1 - 32,400), <i>Minutes</i> (1 - 540) or <i>Hours</i> (1 - 9). The default setting is 30 seconds.
UDP	Define a flow timeout value in either <i>Seconds</i> (1 - 32,400), <i>Minutes</i> (1 - 540) or <i>Hours</i> (1 - 9). The default setting is 90 seconds.
Any Other Flow	Define a flow timeout value in either <i>Seconds</i> (1 - 32,400), <i>Minutes</i> (1 - 540) or <i>Hours</i> (1 - 9). The default setting is 5 seconds.

9. Refer to the **TCP Protocol Checks** field to set the following parameters:

Check TCP states where a SYN packet tears down the flow	Select the check box to allow a SYN packet to delete an old flow in TCP_FIN_FIN_STATE and TCP_CLOSED_STATE and create a new flow. The default setting is enabled.
Check unnecessary resends of TCP packets	Select the check box to enable the checking of unnecessary resends of TCP packets. The default setting is enabled.
Check Sequence Number in ICMP Unreachable error packets	Select the check box to enable sequence number checks in ICMP unreachable error packets when an established TCP flow is aborted. The default setting is enabled.
Check Acknowledgment Number in RST packets	Select the check box to enable the checking of the acknowledgment number in RST packets which aborts a TCP flow in the SYN state. The default setting is enabled.
Check Sequence Number in RST packets	Select the check box to check the sequence number in RST packets which abort an established TCP flow. The default setting is enabled.

10. Select **OK** to update the Firewall Policy Advanced Settings. Select **Reset** to revert to the last saved configuration.

Configuring IP Firewall Rules

Wireless Firewall

Controllers use IP based Firewalls like *Access Control Lists* (ACLs) to filter/mark packets based on the IP from which they arrive, as opposed to filtering packets on Layer 2 ports.

IP based Firewall rules are specific to source and destination IP addresses and the unique rules and precedence orders assigned. Both IP and non-IP traffic on the same Layer 2 interface can be filtered by applying an IP ACL.

NOTE

Once defined, a set of IP Firewall rules must be applied to an interface to be a functional filtering tool.

To add or edit an IP based Firewall Rule policy:

1. Select **Configuration > Security > Wireless Firewall > IP Firewall Rules** to display existing IP Firewall Rule policies.




FIGURE 294 IP Firewall Rules screen

2. Select **+ Add Row** to create a new IP Firewall Rule. Select an existing policy and click **Edit** to modify the attributes of that rule configuration.
3. Select the added row to expand it into configurable parameters for defining the IP based Firewall rule.

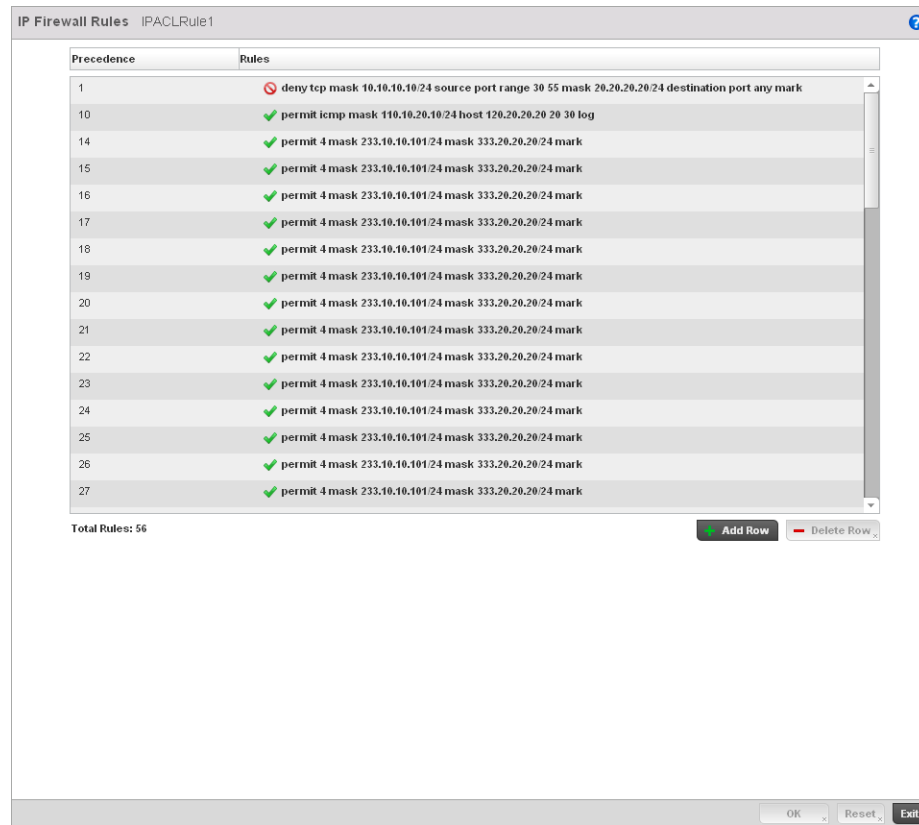


FIGURE 295 IP Firewall Rules Add screen

4. If adding a new **IP Firewall Rule**, provide a name up to 32 characters.
5. Define the following parameters for the IP Firewall Rule:

Allow	Every IP Firewall rule is made up of matching criteria rules. The action defines what to do with the packet if it matches the specified criteria. The following actions are supported: <i>Deny</i> — Instructs the Firewall to not to allow a packet to proceed to its destination. <i>Permit</i> —Instructs the Firewall to allow a packet to proceed to its destination.
Source / Destination	Enter both <i>Source</i> and <i>Destination</i> IP addresses. The wireless controller uses the source IP address, destination IP address and IP protocol type as basic matching criteria. The wireless controller's access policy filter can also include other parameters specific to a protocol type (like source and destination port for TCP/UDP protocol. Provide a subnet mask if needed.
Protocol	Select the protocol used with the IP rule from the drop-down menu. IP is selected by default. Selecting ICMP displays an additional set of ICMP specific options for ICMP type and code. Selecting either TCP or UDP displays an additional set of specific TCP/UDP source and destinations port options.

Action	The following actions are supported: <i>Log</i> —Events are logged to the controller for archive and analysis. <i>Mark</i> —Modifies certain fields inside the packet and then permits them. Therefore, mark is an action with an implicit permit. - VLAN 802.1p priority. - DSCP bits in the IP header. - TOS bits in the IP header. <i>Mark, Log</i> — Conducts both mark and log functions.
Precedence	Use the spinner control to specify a precedence for this IP policy between 1-1500. Rules with lower precedence are always applied first to packets.
Description	Provide a description up to characters long for rule to help differentiate it from others with similar configurations.

6. Select **+ Add Row** as needed to add additional IP Firewall Rule configurations. Select the **Delete Row** icon as required to remove selected IP Firewall Rules..
7. Select **OK** when completed to update the IP Firewall rules. Select **Reset** to revert the screen back to its last saved configuration.

Configuring MAC Firewall Rules

Wireless Firewall

Controllers can use MAC based Firewalls like *Access Control Lists (ACLs)* to filter/mark packets based on the IP from which they arrive, as opposed to filtering packets on Layer 2 ports.

Optionally filter Layer 2 traffic on a physical Layer 2 interface using MAC addresses. A MAC Firewall rule uses source and destination MAC addresses for matching operations, where the result is a typical allow, deny or mark designation to RF Domain manager packet traffic.

NOTE

Once defined, a set of MAC Firewall rules must be applied to an interface to be a functional filtering tool.

To add or edit a MAC based Firewall Rule policy:

1. Select **Configuration > Security > Wireless Firewall > MAC Firewall Rules** to display existing IP Firewall Rule policies.




FIGURE 296 MAC Firewall Rules screen

2. Select **+ Add Row** to create a new MAC Firewall Rule. Select an existing policy and click **Edit** to modify the attributes of that rule's configuration.
3. Select the added row to expand it into configurable parameters for defining the MAC based Firewall rule.

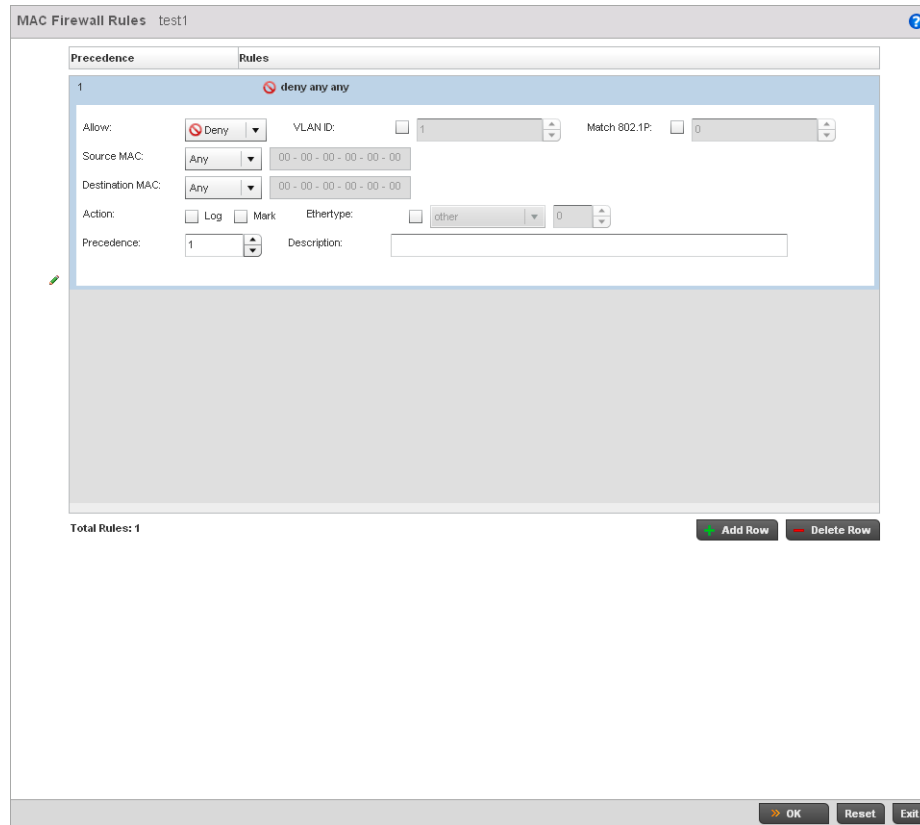


FIGURE 297 MAC Firewall Rules screen

4. If adding a new **MAC Firewall Rule**, provide a name up to 32 characters.
5. Define the following parameters for the IP Firewall Rule:

Allow	Every IP Firewall rule is made up of matching criteria rules. The action defines what to do with the packet if it matches the specified criteria. The following actions are supported: <i>Deny</i> — Instructs the Firewall to prevent a packet from proceeding to its destination. <i>Permit</i> —Instructs the Firewall to allow a packet to proceed to its destination.
VLAN ID	Enter a VLAN ID representative of the shared SSID each user employs to interoperate within the managed network (once authenticated by the local RADIUS server). The VLAN ID can be between 1 and 4094.
Match 802.1P	Configures IP DSCP to 802.1p priority mapping for untagged frames. Use the spinner control to define a setting between 0-7.
Source and Destination MAC	Enter both <i>Source</i> and <i>Destination</i> MAC addresses. The wireless controller uses the source IP address, destination MAC address as basic matching criteria. Provide a subnet mask if using a mask.

Action	<p>The following actions are supported:</p> <p><i>Log</i>—Events are logged to the controller for archive and analysis.</p> <p><i>Mark</i>—Modifies certain fields inside the packet and then permits them. Therefore, mark is an action with an implicit permit.</p> <ul style="list-style-type: none"> - VLAN 802.1p priority. - DSCP bits in the IP header. - TOS bits in the IP header. <p><i>Mark, Log</i> — Conducts both mark and log functions.</p>
Ethertype	<p>Use the drop-down menu to specify an Ethertype of either <i>ipv6</i>, <i>arp</i>, <i>wisp</i>, or <i>monitor 8021q</i>. An EtherType is a two-octet field within an Ethernet frame. It's used to indicate which protocol is encapsulated in the payload of an Ethernet frame.</p>
Precedence	<p>Use the spinner control to specify a precedence for this MAC Firewall rule between 1-1500. Rules with lower precedence are always applied first to packets.</p>
Description	<p>Provide a description (up to 64 characters) for the rule to help differentiate the it from others with similar configurations.</p>

6. Select **+ Add Row** as needed to add additional MAC Firewall Rule configurations. Select the **Delete Row** icon as required to remove selected MAC Firewall Rules.
7. Select **OK** when completed to update the MAC Firewall Rules. Select **Reset** to revert the screen back to its last saved configuration.

Firewall Deployment Considerations

[Configuring a Firewall Policy](#)

Before defining a Firewall configuration on the wireless controller, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- Firewalls implement access control policies, so if you don't have an idea of what kind of access to allow or deny, a Firewall is of little value.
- It's important to recognize the Firewall's configuration is a mechanism for enforcing a wireless controller managed network access policy.
- A role based Firewall requires an advanced security license to apply inbound and outbound Firewall policies to users and devices
- Firewalls cannot protect against tunneling over application protocols to poorly secured wireless clients.
- Firewalls should be deployed on WLANs implementing weak encryption to minimize access to trusted networks and hosts in the event the WLAN is compromised.
- Firewalls should be enabled when providing managed Hotspot guest access. Firewall policies should be applied to Hotspot enabled WLANs to prevent guest user traffic from being routed to trusted networks and hosts.

Wireless Client Roles

Define wireless client roles to filter clients from controller interoperation based on matching policies. Matching policies (much like ACLs) are sequential collections of permit and deny conditions that apply to packets received from controller connected clients. When a packet is received from a client, the controller compares the fields in the packet against applied matching policy rules to verify the packet has the required permissions to be forwarded, based on the criteria specified. If a packet does not meet any of the criteria specified, the packet is dropped.

Additionally, wireless client connections are also managed by granting or restricting access by specifying a range of IP or MAC addresses to include or exclude from controller connectivity. These MAC or IP access control mechanisms are configured as Firewall Rules to further refine client filter and matching criteria.

Configuring a Client's Role Policy

[Wireless Client Roles](#)

To configure a wireless client's role policy and matching criteria:

1. Select **Configuration > Security > Wireless Client Roles**.

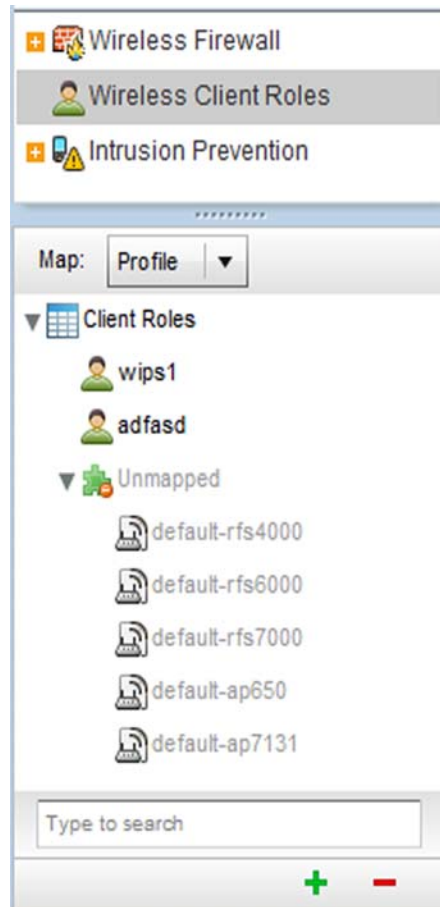


FIGURE 298 Configuration > Security screen and Client Roles Browser

2. The **Wireless Client Roles** screen displays the name of those client role policies created thus far. Either select **Add** to create a new Wireless Client Role policy, **Edit** to modify an existing policy or **Delete** to remove a policy.




FIGURE 299 Wireless IPS screen

3. The **Roles** tab displays by default. If no policies have been created, a default wireless client role policy can be applied. The Roles screen lists existing policies. Any of these existing policies can be selected and applied.




FIGURE 300 Wireless IPS Roles screen

4. Refer to the following configuration data for existing Wireless IPS policies:

Role Name	Displays the name assigned to the client role policy when it was initially created.
Precedence	Displays the precedence number associated with each role. Precedence numbers determine the order a role is applied. Roles with lower numbers are applied before those with higher numbers. Precedence numbers are assigned when a role is created or modified, and two or more roles can share the same precedence.

5. Select **Add** to create a new wireless client role policy, **Edit** to modify the attributes of a selected policy or **Delete** to remove obsolete policies from the list of those available to the controller.
6. The Role Policy Roles screen displays with the **Settings** tab displayed by default.

FIGURE 301 Wireless Client Roles screen - Settings tab

7. If creating a new role, assign it name to help differentiate it from others that may have a similar configuration. The role policy name cannot exceed 64 characters. The name cannot be modified as part of the edit process.
8. Within the **Role Precedence** field, use the spinner control to set a numerical precedence value between 1 - 10,000. Precedence determines the order a role is applied. Roles with lower numbers are applied before those with higher numbers. While there's no default precedence for a role, two or more roles can share the same precedence.

9. Refer to the **Match Expressions** field to create filter rules based on AP locations, SSIDs and RADIUS group memberships.

AP Location	<p>Use the drop-down menu to specify the location of an Access Point matched in a RF Domain or the Access Point's resident configuration. Select one of the following filter options:</p> <p><i>Exact</i> - The role is only applied to APs with the exact location string specified in the role.</p> <p><i>Contains</i> - The role is only applied to APs whose location contains the location string specified in the role.</p> <p><i>Does Not Contain</i> - The role is only applied to APs whose location does not contain the location string specified in the role.</p> <p><i>Any</i> - The role is applied to any AP location. This is the default setting.</p>
SSID Configuration	<p>Use the drop-down menu to define a wireless client filter option based on how the SSID is specified in a WLAN. Select one of the following options:</p> <p><i>Exact</i> -The role is only applied when the exact SSID string specified in the role.</p> <p><i>Contains</i> -The role is only applied when the SSID contains the string specified in the role.</p> <p><i>Does Not Contain</i> - The role is applied when the SSID does not contain the string specified in the role.</p> <p><i>Any</i> - The role is applied to any SSID Location. This is the default setting.</p>
Group Configuration	<p>Use the drop-down menu to define a wireless client filter option based on how the RADIUS group name matches the provided expression. Select one of the following options:</p> <p><i>Exact</i> -The role is only applied when the exact Radius Group Name string is specified in the role.</p> <p><i>Contains</i> - The role is applied when the Radius Group Name contains the string specified in the role.</p> <p><i>Does Not Contain</i> - The role is applied when the Radius Group Name does not contain the string specified in the role</p> <p><i>Any</i> - The role is applied to any RADIUS group name. This is the default setting.</p>

10. Use the **Wireless Client Filter** parameter to define a wireless client MAC address filter that's applied to each role. Select the **Any** radio button to use any MAC address. The default setting is **Any**.
11. Refer to the **Captive Portal Connection** parameter to define when wireless clients are authenticated when making a captive portal authentication request to the controller.

Secure guest access to the controller is referred to as *captive portal*. A captive portal is guest access policy for providing temporary and restrictive access to the managed wireless network. The primary means of securing guest access is a hotspot. Existing captive portal policies can be applied to a WLAN to provide secure guest access.
12. Select the **Pre-Login** check box to conduct captive portal client authentication before the client is logged into the controller. Select **Post-Login** to have the client share authentication credentials with the controller after it has logged into the managed network. Select **Any** (the default setting) makes no distinction on whether authentication is conducted before or after the client has logged in.
13. Use the **Authentication / Encryption** field to set the authentication and encryption filters applied to this wireless client role. The options for both authentication and encryption are:
 - *Equals* - The role is only applied when the authentication and encryption type matches the exact method(s) specified by the radio button selections.

- *Not Equals* - The role is only applied when the authentication and encryption type does not match the exact method(s) specified by the radio button selections.
 - *Any* - The role is applied to any type. This is the default setting for both authentication and encryption.
14. Select **OK** to update the Settings screen. Select **Reset** to revert to the last saved configuration.
 15. Select the **Firewall Rules** tab to set default Firewall rules for *Inbound* and *Outbound* IP and MAC Firewall rules.

The screenshot displays the 'Firewall Rules' tab for a role named 'test1'. It is divided into four quadrants: IP Inbound, MAC Inbound, IP Outbound, and MAC Outbound. Each quadrant features a table with two columns: 'Firewall Rules Name' and 'Precedence', along with a delete icon. An 'Add Row' button is positioned below each table. At the bottom of the screen, there are three buttons: 'OK', 'Reset', and 'Exit'.

FIGURE 302 Wireless Client Roles screen - Default Firewall Rules tab

A Firewall is a mechanism enforcing access control, and is considered a first line of defense in protecting proprietary information within the network. The means by which this is accomplished varies, but in principle, a Firewall can be thought of as mechanisms both *blocking* and *permitting* data traffic based on inbound and outbound IP and MAC rules.

IP based Firewall rules are specific to source and destination IP addresses and the unique rules and precedence orders assigned. Both IP and non-IP traffic on the same Layer 2 interface can be filtered by applying both an IP ACL and a MAC.

Additionally, the controller allows administrators to filter Layer 2 traffic on a physical Layer 2 interface using MAC addresses. A MAC Firewall rule uses source and destination MAC addresses for matching operations, where the result is a typical allow, deny or mark designation to controller packet traffic.

16. Specify an **IP Inbound** or **IP Outbound** Firewall rule by selecting a rule from the drop-down menu and use the spinner control to assign the rule Precedence. Rules with lower precedence are always applied first to packets.
17. If no IP Inbound or Outbound rules exist meeting the required Firewall filtering criteria, select the **Create** button to set the inbound or outbound rule criteria. Select the **+ Add Row** button or **Delete** icon as needed to add or remove IP Firewall rules. Define the following parameters to create a new Inbound or Outbound IP Firewall rule:

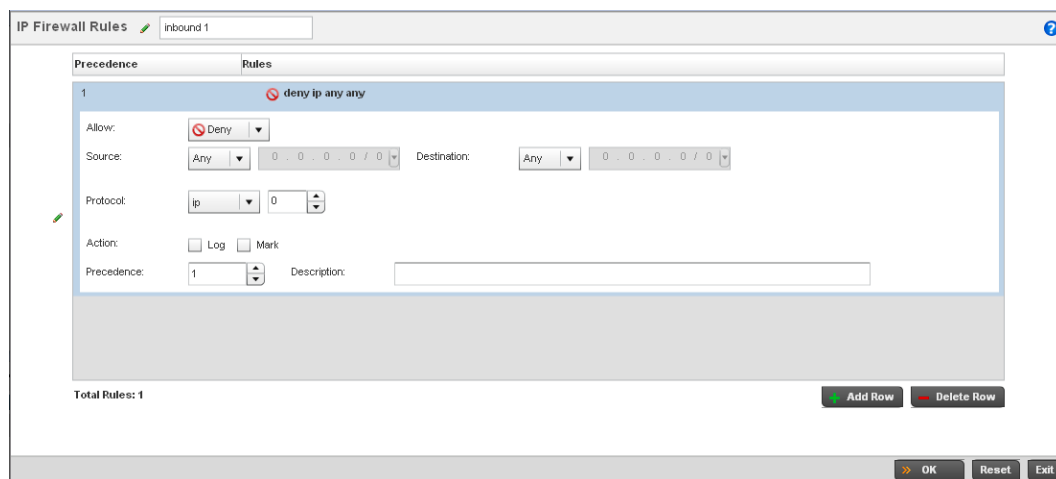


FIGURE 303 Wireless Client Roles - IP Firewall Rules screen

IP Firewall Rules

If creating a new IP Firewall rule, assign it a name (up to 64 characters) to help differentiate it from others that may have similar configurations.

Allow

Every IP Firewall rule is made up of matching criteria rules. The action defines what to do with the packet if it matches the specified criteria. The following actions are supported:

Deny—Instructs the Firewall to prohibit a packet from proceeding to its destination.

Permit—Instructs the Firewall to allow a packet to proceed to its destination.

Source / Destination

Enter both Source and Destination IP addresses. The wireless controller uses the source IP address, destination IP address and IP protocol type as basic matching criteria. The wireless controller's access policy filter can also include other parameters specific to a protocol type (like source and destination port for TCP/UDP protocols).

Protocol

Select the *IP*, *ICMP*, *TCP* or *UDP* protocol used with the IP access policy. IP is selected by default. Selecting ICMP displays an additional set of ICMP specific options to set the ICMP *Type* and *Code*. Selecting either TCP or UDP displays an additional set of specific TCP/UDP source and destinations port options.

Action

The following actions are supported:

Log—Logs the event when this rule is applied to a wireless clients association attempt.

Mark—Modifies certain fields inside the packet and then permits them.

Therefore, mark is an action with an implicit permit.

- VLAN 802.1p priority.

- DSCP bits in the IP header.

- TOS bits in the IP header.

Mark, Log — Applies both log and mark actions.

Precedence

Use the spinner control to specify a precedence for this IP policy between 1-1500. Rules with lower precedence are always applied first. More than one rule can share the same precedence value.

Description

Provide a description of the rule to differentiate it from others with similar configurations. This should be more descriptive than simply re-applying the name of the rule.

18. Select **OK** to save the updates to the Inbound or Outbound IP Firewall rule. Select **Reset** to revert to the last saved configuration.
19. If required, select existing Inbound and Outbound MAC Firewall Rules using the drop-down menu. If no rules exist, select **Create** to display a screen where Inbound or Outbound Firewall rules can be created.
20. Define the following parameters required to create an Inbound or Outbound MAC Firewall rule:

FIGURE 304 Wireless Client Roles - MAC Firewall Rules screen

MAC Firewall Rules	If creating a new MAC Firewall rule, assign it a name (up to 64 characters) to help differentiate it from others that may have similar configurations.
Allow	Every MAC Firewall rule is made up of matching criteria rules. The action defines what to do with the packet if it matches the specified criteria. The following actions are supported: <i>Deny</i> —Instructs the Firewall to not to allow a packet to proceed to its destination. <i>Permit</i> —Instructs the Firewall to allow a packet to proceed to its destination.
VLAN ID	Enter a VLAN ID representative of the shared SSID each user employs to interoperate within the network (once authenticated by the local RADIUS server). The VLAN ID can be between 1 and 4094.
Match 802.1P	Configures IP DSCP to 802.1p priority mapping for untagged frames. Use the spinner control to define a setting between 0-7.
Source / Destination MAC	Enter both <i>Source</i> and <i>Destination</i> MAC addresses. The wireless controller uses the addresses as basic matching criteria.
Action	The following actions are supported: <i>Log</i> —Logs the event when this rule is applied to a wireless clients association attempt. <i>Mark</i> —Modifies certain fields inside the packet and then permits them. Therefore, mark is an action with an implicit permit. - VLAN 802.1p priority. - DSCP bits in the header. - TOS bits in the header. <i>Mark, Log</i> — Applies both log and mark actions.

Ethertype	Use the drop-down menu to specify an Ether type. An EtherType is a two-octet field within an Ethernet frame. It's used to indicate which protocol is encapsulated in the payload of an Ethernet frame.
Precedence	Use the spinner control to specify a precedence for this MAC policy between 1-1500. Rules with lower precedence are always applied first to packets. More than one rule can share the same precedence value.
Description	Provide a description for the rule to differentiate the IP Firewall Rule from others with similar configurations. This should be more descriptive than simply re-applying the name of the rule.

21. Select **OK** to save the updates to the MAC Firewall rule. Select **Reset** to revert to the last saved configuration.

Intrusion Prevention

The wireless controller supports *Wireless Intrusion Protection Systems (WIPS)* to provide continuous protection against wireless threats and act as an additional layer of security complementing wireless VPNs and encryption and authentication policies. The wireless controller supports WIPS through the use of dedicated sensor devices designed to actively detect and locate unauthorized AP devices. After detection, they use mitigation techniques to block the devices by manual termination or air lockdown.

Unauthorized APs are untrusted access points connected to a wireless controller managed LAN that accept client associations. They can be deployed for illegal wireless access to a corporate network, implanted with malicious intent by an attacker, or could just be misconfigured access points that do not adhere to corporate policies. An attacker can install a unauthorized AP with the same ESSID as the authorized WLAN, causing a nearby client to associate to it. The unauthorized AP can then steal user credentials from the client, launch a man-in-the middle attack or take control of wireless clients to launch denial-of-service attacks.

Brocade Mobility RFS4000, RFS6000 and RFS7000 wireless controllers support unauthorized AP detection, location and containment natively. A WIPS server can alternatively be deployed (in conjunction with the wireless controller) as a dedicated solution within a separate enclosure. When used within a Brocade Mobility wireless controller managed network and its associated access point radios, a WIPS deployment provides the following enterprise class security management features and functionality:

- *Threat Detection* - Threat detection is central to a wireless security solution. Threat detection must be robust enough to correctly detect threats and swiftly help protect the wireless controller managed wireless network.
- *Rogue Detection and Segregation* - A WIPS supported wireless controller distinguishes itself by both identifying and categorizing nearby APs. WIPS identifies threatening versus non-threatening APs by segregating APs attached to the network (unauthorized APs) from those not attached to the network (neighboring APs). The correct classification of potential threats is critical in order for administrators to act promptly against rogues and not invest in a manual search of neighboring APs to isolate the few attached to the network.
- *Locationing* - Administrators can define the location of wireless clients as they move throughout a wireless controller managed site. This allows for the removal of potential rogues though the identification and removal of their connected access points.

- *WEP Cloaking* - WEP Cloaking protects organizations using the *Wired Equivalent Privacy* (WEP) security standard to protect networks from common attempts used to crack encryption keys. There are several freeware WEP cracking tools available and 23 known attacks against the original 802.11 encryption standard; even 128-bit WEP keys take only minutes to crack. WEP Cloaking module enables organizations to operate WEP encrypted networks securely and to preserve their existing investment in mobile devices.

Configuring a WIPS Policy

Intrusion Prevention

To configure WIPS on the wireless controller:

1. Select **Configuration > Security > Intrusion Prevention**.
2. Expand the Intrusion Prevention option within the Configuration > Security menu to display the *WIPS Policy*, *Advanced WIPS Policy* and *Device Categorization* items available.

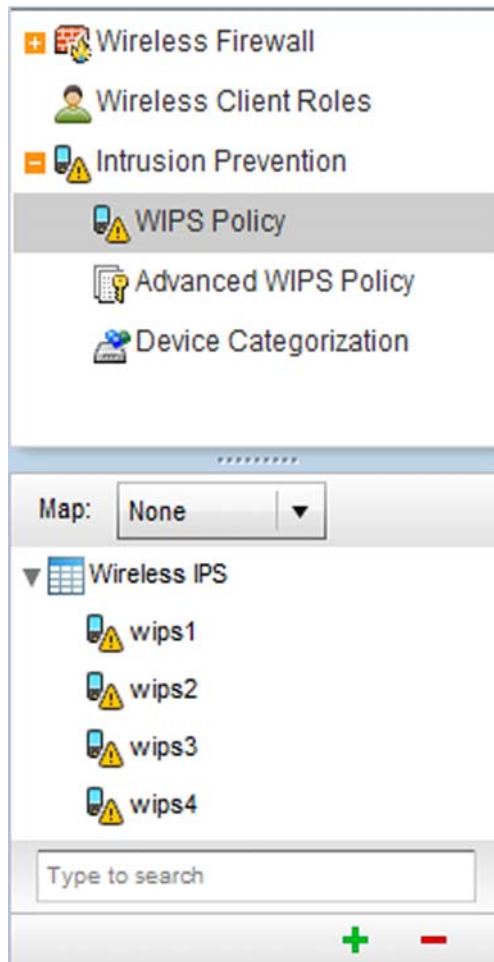
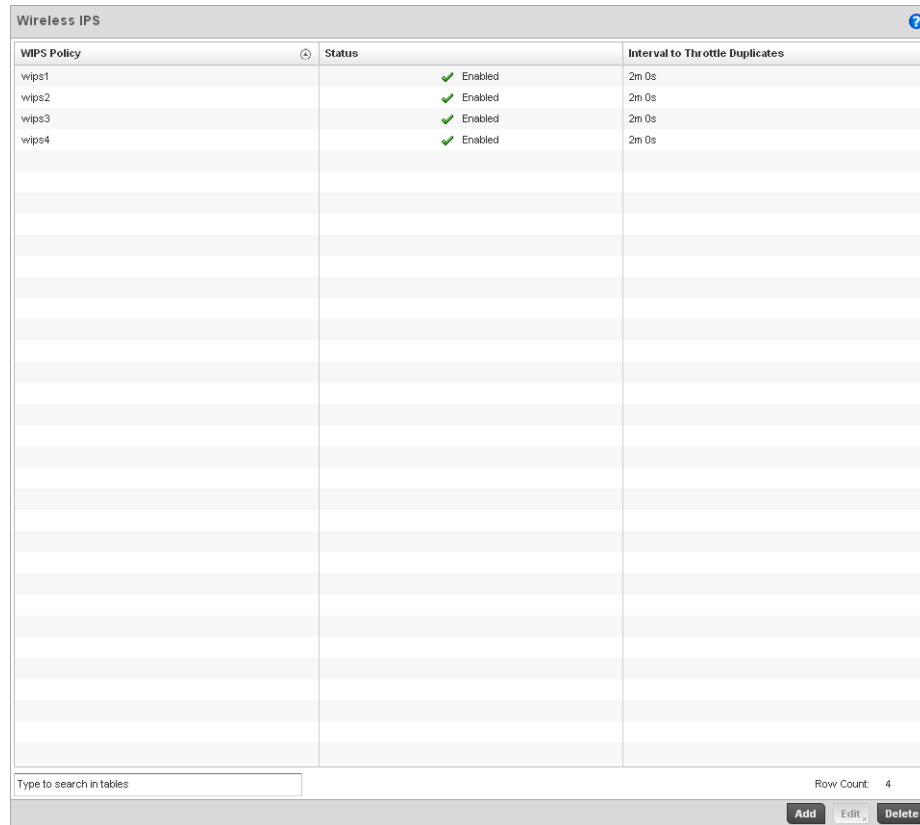


FIGURE 305 Configuration > Security screen

3. The Wireless IPS screen displays by default. The Wireless IPS screen lists existing WIPS policies if any are configured. Any of these existing WIPS policies can be selected and applied.



WIPS Policy	Status	Interval to Throttle Duplicates
wips1	Enabled	2m 0s
wips2	Enabled	2m 0s
wips3	Enabled	2m 0s
wips4	Enabled	2m 0s

Type to search in tables

Row Count: 4

FIGURE 306 Wireless IPS screen

4. Refer to the following for existing Wireless IPS policies:

WIPS Policy	Displays the name assigned to the WIPS policy when it was initially created. The name cannot be modified as part of the edit process.
Status	Displays a green checkmark if the listed WIPS policy is enabled and ready for use with a controller profile. A red "X" designated the listed WIPS policy as disabled.
Interval to Throttle Duplicates	Displays the duration in seconds when traffic meeting the criteria defined in the selected WIPS policy is prevented/throttled.

5. Select **Add** to create a new WIPS policy, **Edit** to modify the attributes of a selected policy or **Delete** to remove obsolete policies from the list of those available to the controller.
6. If adding or editing an existing WIPS policy, the WIPS Policy screen displays with the settings tab displayed by default.

FIGURE 307 WIPS Policy screen - Settings tab

7. If creating a new **WIPS Policy**, assign it name to help differentiate it from others that may have a similar configuration. The policy name cannot exceed 64 characters. The name cannot be modified as part of the edit process.
8. Within the **Wireless IPS Status** field, select either the *Enabled* or *Disabled* radio button to either activate or de-activate the WIPS policy for use with a controller profile. The default setting is disabled.
9. Enter the **Interval to Throttle Packets** in either *Seconds* (1 - 86,400), *Minutes* (1 - 1,400), *Hours* (1 - 24) or *Days* (1). This interval represents the duration event duplicates are *not* stored in history. The default setting is 120 seconds.
10. Refer to the **Rogue AP Detection** field to define the following detection settings for this WIPS policy:

Enable Rogue AP Detection	Select the checkbox to enable the detection of unauthorized (unsanctioned) devices fro this WIPS policy. The default setting is disabled.
Wait Time to Determine AP Status	Define a wait time in either <i>Seconds</i> (10 - 600) or <i>Minutes</i> (1 - 10) before a detected AP is interpreted as a rogue (unsanctioned) device, and potentially removed from controller management. The default interval is 1 minute.
Ageout for AP Entries	Set the interval the WIPS policy uses to ageout rogue devices. Set the policy in either <i>Seconds</i> (30 - 86,400), <i>Minutes</i> (1- 1,440), <i>Hours</i> (1 - 24) or <i>Days</i> (1). The default setting is 5 minutes.

11. Use the **Device Categorization Policy** drop-down menu to select a policy describing whether a device is filtered as sanctioned, a client or Access Point and the MAC and SSID addresses used as filtering mechanisms.
12. If a policy requires creation, select the **Create** button. If an existing policy requires modification, select the **Edit** button and update the Device Categorization Policy as needed.

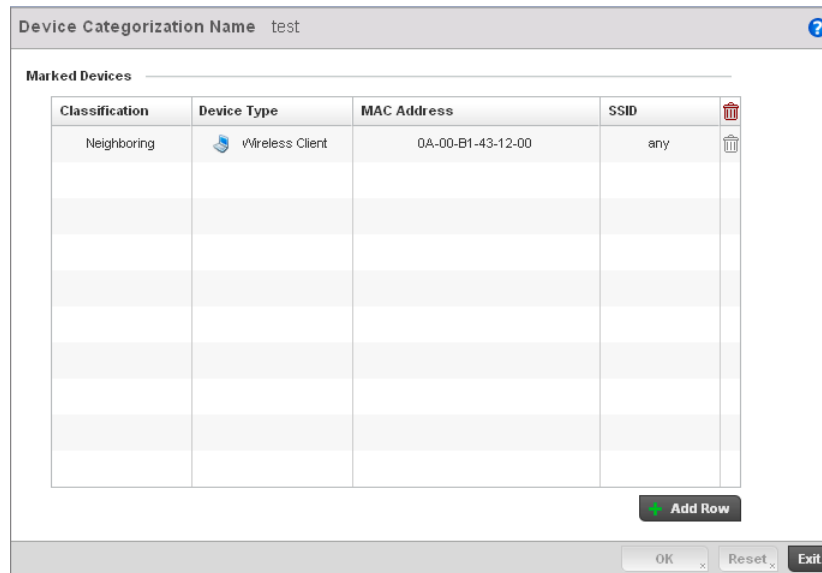


FIGURE 308 Device Categorization screen

13. Select **OK** to update the settings. Select **Reset** to revert to the last saved configuration.
14. Select the **WIPS Events** tab to enable events, filters and threshold values for this WIPS policy. The **Excessive** tab displays by default.

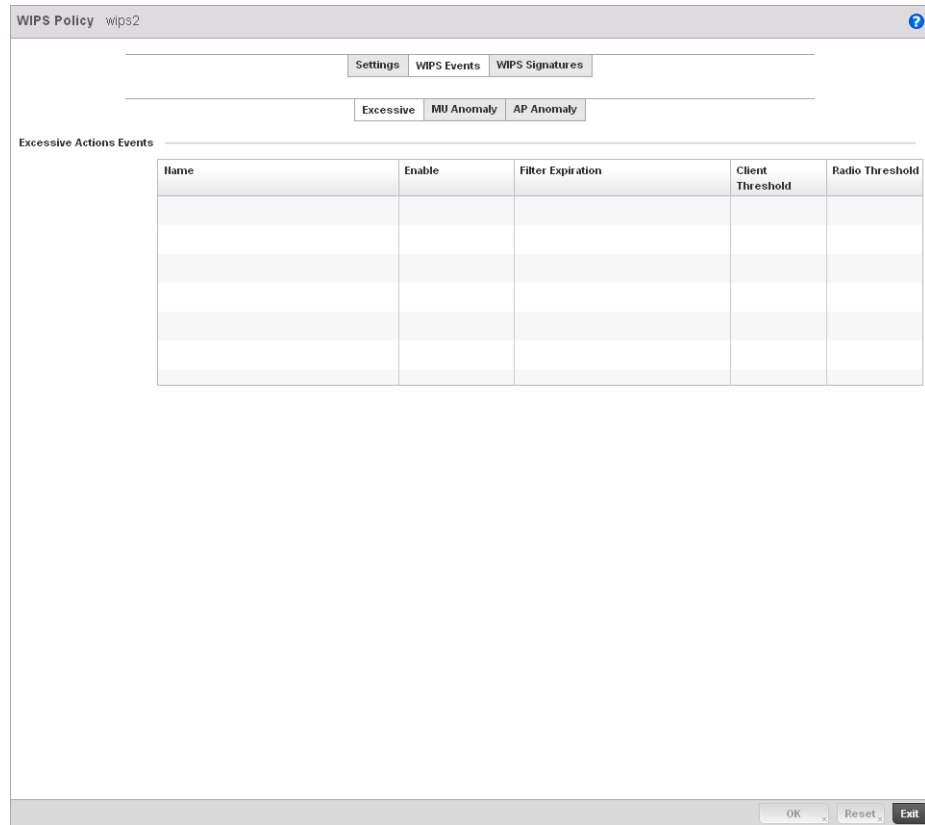


FIGURE 309 WIPS Events screen - Excessive tab

The Excessive tab lists a series of events that can impact the performance of the managed network. An administrator can enable or disable the filtering of each listed event and set the thresholds required for the generation of the event notification and filtering action applied.

An Excessive Action Event is an event where an action is performed repetitively and continuously. DoS attacks come under this category. Use the *Excessive Action Events* table to select and configure the action taken when events are triggered.

15. Set the configurations of the following **Excessive Action Events**:

Name	Displays the name of the excessive action event. This column lists the event being tracked against the defined thresholds set for interpreting the event as excessive or permitted.
Enable	Displays whether tracking is enabled for each Excessive Action Event. Use the drop-down menu to enable/disable events as required. A green checkmark defines the event as enabled for tracking against its threshold values. A red "X" defines the event as disabled and not tracked by the WIPS policy. Each event is disabled by default.

Filter Expiration

Set the duration the anomaly causing client is filtered. This creates a special ACL entry and frames coming from the client are silently dropped. The default setting is 0 seconds.

This value is applicable across the RF Domain. If a station is detected performing an attack and is filtered by one of the APs, the information is passed to the domain controller. The domain controller then propagates this information to all APs and controllers in the RF Domain.

Client Threshold

Set the client threshold after which the filter is triggered and an event generated.

Radio Threshold

Set the radio threshold after which an event is recorded to the events history.

16. Select **OK** to save the updates to the to Excessive Actions configuration used by the WIPS policy. Select **Reset** to revert to the last saved configuration.

17. Select the **MU Anomaly** tab:

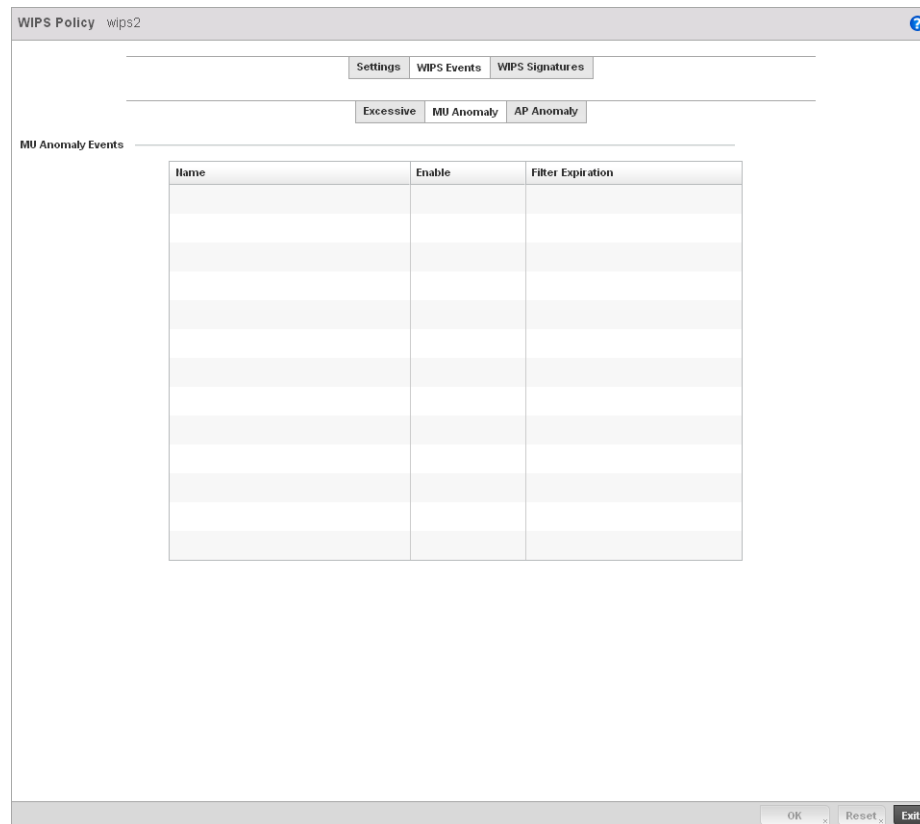


FIGURE 310 WIPS Events screen - MU Anomaly tab

MU Anomaly events are suspicious events performed by wireless clients that can compromise the security and stability of the network. Use this MU Anomaly screen to configure the intervals clients can be filtered upon the generation of each defined event.

18. Set the configurations of the following **MU Anomaly Events** configurations:

- Name** Displays the name of the MU Anomaly event. This column lists the event being tracked against the defined thresholds set for interpreting the event as excessive or permitted.
- Enable** Displays whether tracking is enabled for each MU Anomaly event. Use the drop-down menu to enable/disable events as required. A green checkmark defines the event as enabled for tracking against its threshold values. A red "X" defines the event as disabled and not tracked by the WIPS policy. Each event is disabled by default.
- Filter Expiration** Set the duration the anomaly causing client is filtered. This creates a special ACL entry and frames coming from the client are silently dropped. The default setting is 0 seconds. For each violation, define a time to filter value in seconds which determines how long the controller ignores packets received from an attacking device once a violation has been triggered. Ignoring frames from an attacking device minimizes the effectiveness of the attack and the impact to the site until permanent mitigation can be performed.

19. Select **OK** to save the updates to the MU Anomaly configuration used by the WIPS policy. Select **Reset** to revert to the last saved configuration.

20. Select the **AP Anomaly** tab.

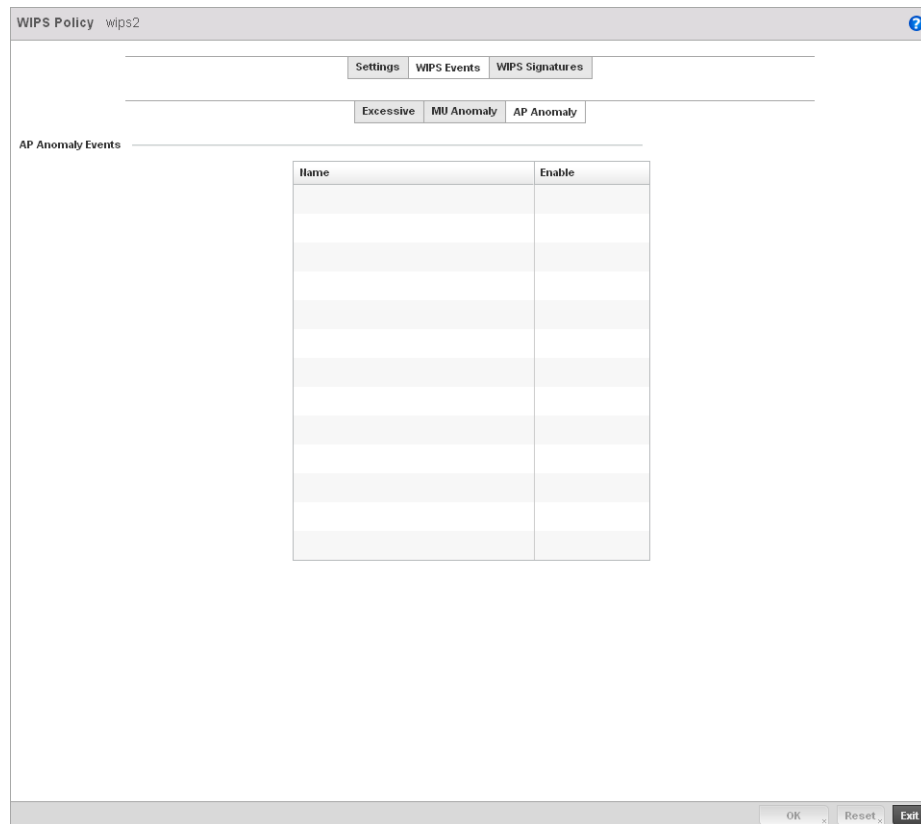


FIGURE 311 WIPS Events screen - AP Anomaly tab

AP Anomaly events are suspicious frames sent by a neighboring APs. Use this screen to determine whether an event is enabled for tracking.

21. Set the configurations of the following MU Anomaly Events configurations:

- Name** Displays the name of the MU Anomaly event. This column lists the event being tracked against the defined thresholds set for interpreting the event as excessive or permitted.
- Enable** Displays whether tracking is enabled for each MU Anomaly event. Use the drop-down menu to enable/disable events as required. A green checkmark defines the event as enabled for tracking against its threshold values. A red "X" defines the event as disabled and not tracked by the WIPS policy. Each event is disabled by default.

22. Select **OK** to save the updates to the AP Anomaly configuration used by the WIPS policy. Select **Reset** to revert to the last saved configuration.

23. Select the **WIPS Signatures** tab.

Name	Signature	BSSID MAC	Source MAC	Destination MAC	Frame Type to Match	Match on SSID
signature3	✓	Not Set	Not Set	Not Set	Beacon	dodddd

FIGURE 312 WIPS Signatures screen

24. The WIPS Signatures tab displays the following read-only data:

- Name** Lists the name assigned to each signature as it was created. A signature name cannot be modified as part of the edit process.
- Signature** Displays whether the signature is enabled. A green checkmark defines the signature as enabled. A red "X" defines the signature as disabled. Each signature is disabled by default.
- BSSID MAC** Displays each BSS ID MAC address used for matching purposes.

- Source MAC** Displays each source MAC address of the packet examined for matching purposes.
- Destination MAC** Displays each destination MAC address of the packet examined for matching purposes.
- Frame Type to Match** Lists the frame types specified for matching with the WIPS signature.
- Match on SSID** Lists each SSID used for matching purposes.

25. Select **Add** to create a new WIPS signature, **Edit** to modify the attributes of a selected WIPS signature or **Delete** to remove obsolete signatures from the list of those available to the controller.

The screenshot shows the 'WIPS Signatures Configuration' screen for a signature named 'sig1'. The interface is divided into several sections:

- Settings:**
 - Enable Signature:
 - BSSID MAC: (00 - 00 - 00 - 00 - 00 - 00)
 - Source MAC: (00 - 00 - 00 - 00 - 00 - 00)
 - Destination MAC: (00 - 00 - 00 - 00 - 00 - 00)
 - Frame Type to Match: All (dropdown)
 - Match on SSID: (text field)
 - SSID Length: 0 (0 to 32)
- Thresholds:**
 - Wireless Client Threshold: 1 (1 to 65,535)
 - Radio Threshold: 1 (1 to 65,535)
- Filter Expiration:**
 - Filter Expiration: 1 (1 to 86,400 seconds)
- Payload:** A table with columns: Index, Pattern, Offset, and a delete icon. Below the table is an 'Add Row' button.

At the bottom of the window are 'OK', 'Reset', and 'Exit' buttons.

FIGURE 313 WIPS Signatures Configuration screen

26. If adding a new WIPS signature, define a **Name** to distinguish it from others with similar configurations. The name cannot exceed 64 characters.

27. Set the following network address information for a new or modified WIPS Signature:

- Enable Signature** Select the check box to enable the WIPS signature for use with the controller profile. The default signature is enabled.
- BSSID MAC** Define a BSS ID MAC address used for matching purposes.
- Source MAC** Define a source MAC address for the packet examined for matching purposes.
- Destination MAC** Set a destination MAC address for the packet examined for matching purposes.
- Frame Type to Match** Use the drop-down menu to select a frame type matching with the WIPS signature.
- Match on SSID** Sets the SSID used for matching. Ensure it's specified properly or the SSID won't be properly filtered.
- SSID Length** Set the character length of the SSID used for matching purposes. The maximum length is 32 characters.

28. Refer to **Thresholds** field to set the thresholds used as filtering criteria.

Client Threshold Specify the threshold limit per client that, when exceeded, signals the event. The configurable range is from 1 - 65,535.

Radio Threshold Specify the threshold limit per radio that, when exceeded, signals the event. The configurable range is from 1 - 65,535.

29. Set a **Filter Expiration** between 1 - 86,400 seconds that specifies the duration a client is excluded from RF Domain manager radio association when responsible for triggering a WIPS event.

30. Refer to the **Payload** table to set a numerical index pattern and offset for the WIPS signature.

31. Select **OK** to save the updates to the WIPS Signature configuration. Select **Reset** to revert to the last saved configuration.

Configuring an Advanced WIPS Policy

Intrusion Prevention

Define an advanced WIPS configuration to optionally remove (terminate) unwanted device connections, and sanction (allow) or unsanction (disallow) specific events within the managed network.

1. Select **Configuration > Security > Intrusion Prevention**.
2. Expand the Intrusion Prevention option within the Configuration > Security menu and select **Advanced WIPS**.

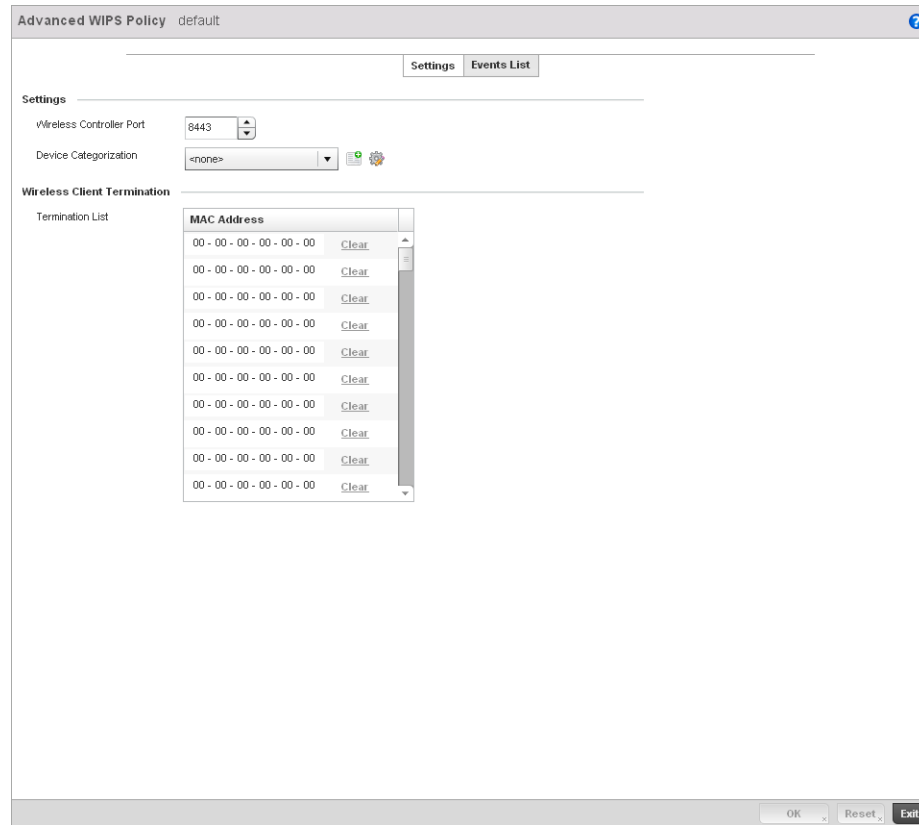


FIGURE 315 Advanced WIPS Policy screen - Settings tab

6. The Advanced WIPS screen displays the **Settings** tab by default.
7. Define the following **Settings** for the Advanced WIPS policy:

Wireless Controller Port Use the spinner control to set the port the advanced WIPS daemon listens over. The default port is 8,443.

Device Categorization Set the device categorization as sanctioned, unsanctioned etc. Select the *Create* icon to create a new Device Category configuration, or the *Edit* icon to modify the configuration of an existing configuration. For information on creating or editing a Device Categorization policy, see *Configuring a WIPS Device Categorization Policy* on page 9-460.

8. Refer to the **Wireless Client Termination** table to set list of up to 100 client MAC that are blocked using this Advanced WIPS policy. This clients are removed from connection within the managed network, so be sure they represent potential threats.
9. Select **OK** to save the updates to the **Advanced WIPS** Settings tab. Select **Reset** to revert to the last saved configuration.
10. Select the **Events List** tab to display a screen where individual events can be enabled as sanctioned or unsanctioned for the managed network.

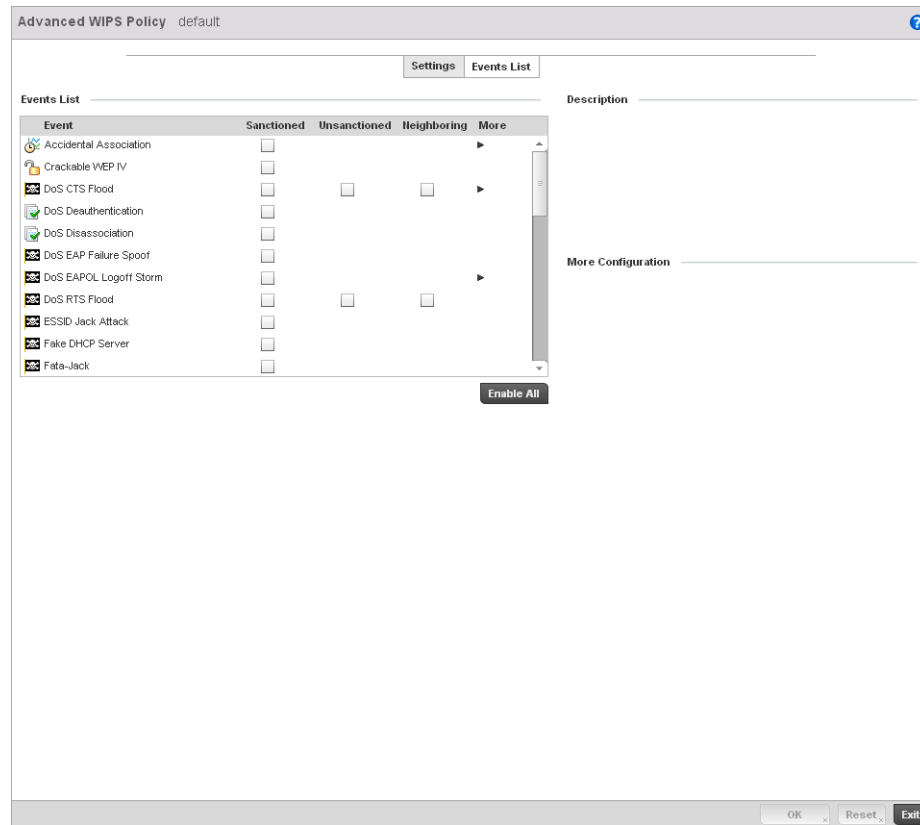


FIGURE 316 Advanced WIPS Policy screen - Events List tab

Events are tracked based on an AP's authorization. APs are either *Sanctioned*, *Unsanctioned* or *Neighboring*. All of the events listed do not necessarily support all the three AP types.

Some events have extra configurable parameters. These events are identified by a small triangle under the **More** column. Extra event parameters are displayed at the right of the screen.

11. Select the radio button corresponding to the **Sanctioned**, **Unsanctioned** or **Neighboring** option for each listed event.
12. Review a description of each event by highlighting it in the table and revising the **Description** displayed on the right-hand of the screen.
13. The **Events List** contains the following events to either authorize, unauthorize or interpret as neighboring for the Advanced WIPS policy:
 - *Accidental Association* - An authorized station has connected to an unauthorized or ignored Access Point.
 - *Crackable WEP IV* - A WEP IV has been detected that could lead to the discovery of the WEP key.
 - *DoS CTS Flood* - An excessive number of CTS frames has been detected.
 - *DoS Deauthentication* - Attack in which deauthentication frames are sent to the wireless client using the MAC address of the AP to which it is associated. This disrupts the client connection and may lead it to associate to a fake AP spoofing the real ESSID.

- *DoS Dissassociation* - A flood of spoofed disassociation frames have been detected.
- *DoS EAP Failure Spoof* - A hacker is sending EAP failure messages to a client using the spoofed MAC address of the Access Point.
- *DoS EAPOL Logoff Storm* - An excessive number of EAPOL Logoff messages has been detected.
- *DoS RTS Flood* - An excessive number of RTS frames has been detected.
- *ESSID Jack Attack* - An active attempt to discover a wireless network's ESSID has been detected.
- *Fake DHCP Server* - A rogue DHCP server is suspected of operating on the wireless network.
- *Fata-Jack* - DoS attack using the Fata-Jack tool, which sends fake authentication failed packets to the wireless client using the spoofed MAC address of the real AP. This leads the client to drop itself from the WLAN.
- *ID Theft EAPOL Success Spoof* - Spoofed EAP success frames have been detected.
- *ID Theft Out-Of-Sequence* - Two devices using the same MAC address have been detected operating in the airspace, resulting in detected wireless frames that are out of sequence.
- *Invalid Channel Advertisement* - An AP is advertising invalid channel.
- *Invalid Management Frame* - Illegal 802.11 management frame has been detected.
- *IPX Detection* - Unencrypted IPX traffic has been observed in the wireless network.
- *Monkey Jack Attack* - Link-layer Man-in-the-Middle attack in which the wireless client associates with a fake access point which then forwards packets between the client and the AP. The attacker may then deny service or perform other attacks on the stream of packets traversing it.
- *NULL Probe Response* - Null probe response frames have been detected with destination of an authorized station.
- *STP Detection* - Unencrypted STP traffic has been observed in the wireless network.
- *Unsanctioned AP* - Unauthorized activity includes events for devices participating in unauthorized communication in your airspace.
- *Windows Zero Config Memory Leak* - Windows XP system memory leak has been detected.
- *WLAN Jack Attack* - DoS attack in which the WLAN Jack tool is used to send de-authentication frames to wireless clients using the spoofed MAC address of the real AP. This leads the clients to de-authenticate and drop their wireless connections.

14. Select **OK** to save the updates to the Advanced WIPS Events List. Select **Reset** to revert to the last saved configuration.

Configuring a WIPS Device Categorization Policy

Intrusion Prevention

Having devices properly classified can help suppress unnecessary unsanctioned AP alarms and allow an administrator to focus on the alarms and devices actually behaving in a suspicious manner. An intruder with a device erroneously authorized could potentially perform activities that harm your organization while appearing to be legitimate. The controller enables devices to be categorized as access points, then defined as sanctioned or unsanctioned for operation within the managed network.

Sanctioned Access Points are generally known to you and conform with your organization's security policies. Unsanctioned devices have been detected as interoperating within the managed network, but are not approved. These devices should be filtered to avoid jeopardizing the data managed by the controller.

To categorize Access Points as sanctioned or unsanctioned within the managed network:

1. Select **Configuration > Security > Intrusion Prevention**.
2. Expand the Intrusion Prevention option within the Configuration > Security menu and select **Device Categorization**.




FIGURE 317 WIPS Device Categorization screen

3. The Device Categorization screen lists those device authorization policies that have been defined thus far.
4. Select **Add** to create a new Device Categorization policy, **Edit** to modify the attributes of a selected existing policy or **Delete** to remove obsolete policies from the list of those available to the controller.

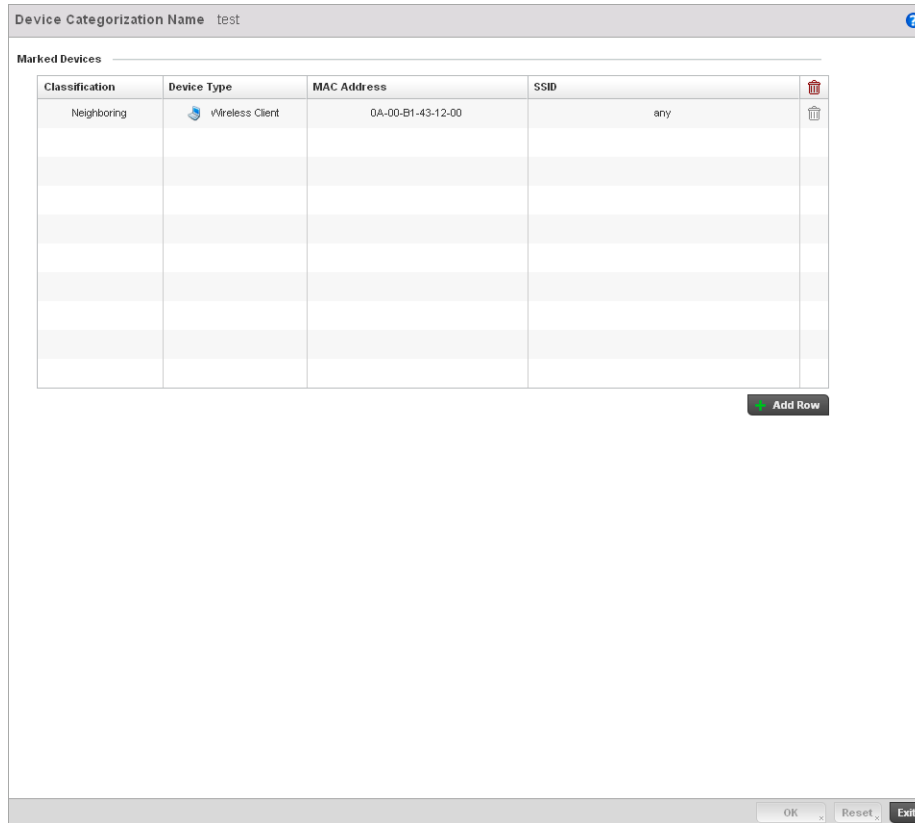


FIGURE 318 WIPS Device Categorization Configuration screen

5. If creating a new Device Categorization policy, provide it a **Name** (up to 64 characters) to distinguish this policy from others with similar configurations. Select **OK** to save the name and enable the remaining parameters on the screen.
6. Select **+ Add Row** to populate the **Marked Devices** field with parameters for adding an Access Point's MAC address, SSID, Access Point designation and controller network authorization. Select the red (-) **Delete Row** icon as needed to remove an individual table entry.
7. Define the following parameters to add a device to a list of devices categorized as sanctioned or unsanctioned for controller network operation:

Classification

Use the drop-down menu to designate the target device as either sanctioned (*True*) or unsanctioned (*False*). The default setting is *False*, categorizing this device as unsanctioned. Thus, each added device requires authorization. A green checkmark designates the device as sanctioned, while a red "X" defines the device as unsanctioned.

Device Type	Use the drop-down menu to designate the target device as either an Access Point (<i>True</i>) or other (<i>False</i>). The default setting is <i>False</i> , categorizing this device as other than an Access Point. A green checkmark designates the device as an Access Point, while a red "X" defines the categorized device as other than an Access Point.
MAC Address	Enter the factory coded MAC address of the target device. This address is hard coded by the device manufacturer and cannot be modified. The MAC address will be defined as sanctioned or unsanctioned as part of the device categorization process.
SSID	Enter the SSID of the target device requiring categorization. The SSID cannot exceed 32 characters.

8. Select **OK** to save the updates to the **Marked Devices** List. Select **Reset** to revert to the last saved configuration.

Intrusion Detection Deployment Considerations

Before configuring WIPS support on the wireless controller, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- WIPS is best utilized when deployed in conjunction with a corporate or enterprise wireless security policy. Since an organization's security goals vary, the security policy should document site specific concerns. The WIPS system can then be modified to support and enforce these additional security policies
- WIPS reporting tools can minimize dedicated administration time. Vulnerability and activity reports should automatically run and be distributed to the appropriate administrators. These reports should highlight areas to be investigated and minimize the need for network monitoring.
- It's important to keep your WIPS system Firmware and Software up to date. A quarterly system audit can ensure firmware and software versions are current.
- Only a trained wireless network administrator can determine the criteria used to authorize or ignore devices. You may want to consider your organization's overall security policy and your tolerance for risk versus users' need for network access. Some questions that may be useful in deciding how to classify a device are:
 - Does the device conform to any vendor requirements you have?
 - What is the signal strength of the device? Is it likely the device is outside your physical radio coverage area?
 - Is the detected Access Point properly configured according to your organization's security policies?
- Brocade Mobility recommends a controller have visibility to all the VLANs deployed. If an external L3 device has been deployed for routing services, each VLAN should be 802.1Q tagged to the controller to allow the controller to detect any unsanctioned APs physically connected to the network.
- Brocade Mobility recommends trusted and known Access Points be added to an sanctioned AP list. This will minimize the number of unsanctioned AP alarms received.

Services Configuration

In this chapter

- [Configuring Captive Portal Policies](#) 465
- [Setting the Controller's DHCP Configuration](#) 477
- [Setting the Controller's RADIUS Configuration](#) 491

The controller natively supports services to provide guest user access to the controller managed network, lease DHCP IP addresses to requesting clients and provide RADIUS client authentication.

Configuring Captive Portal Policies

A *captive portal* is guest access policy for providing guests temporary and restrictive access to the controller managed wireless network. The primary means of securing such controller guest access is the use of a hotspot.

A captive portal policy's hotspot configuration provides secure authenticated controller access using a standard Web browser. Hotspots provides authenticated access by capturing and re-directing a wireless user's Web browser session to a captive portal login page where the user must enter valid credentials to access to the controller managed wireless network. Once logged into the controller managed hotspot, additional Agreement, Welcome and Fail pages provide the administrator with a number of options on the hotspot's screen flow and user appearance.

Hotspot authentication is used primarily for guest or visitor access to the controller managed network, but is increasingly being used to provide authenticated access to private network resources when 802.1X EAP is not a viable option. Hotspot authentication does not provide end-user data encryption, but it can be used with static WEP, WPA-PSK or WPA2-PSK encryption.

Authentication for captive portal guest access applications is performed using a username and password, authenticated by the controller's integrated RADIUS server. Authentication for private network access is provided either locally on the requesting wireless client, or centrally at a datacenter.

The controller uses a Web provisioning tool to create guest user captive portal accounts directly on the controller. The connection medium defined for the controller's Web connection is either HTTP or HTTPS. Both HTTP and HTTPS use a request and response procedure clients follow to disseminate information to and from the controller and requesting wireless clients.

Configuring a Captive Portal Policy

To configure a controller's guest access captive portal policy:

1. Select **Configuration > Services**.

The upper, left-hand, side of the user interface displays a **Services** menu pane where Captive Portal, DHCP and RADIUS configuration options can be selected.

2. Select **Captive Portals**.

The Captive Portal screen displays the configurations of existing policies. New policies can be created, existing policies can be modified or existing policies deleted.

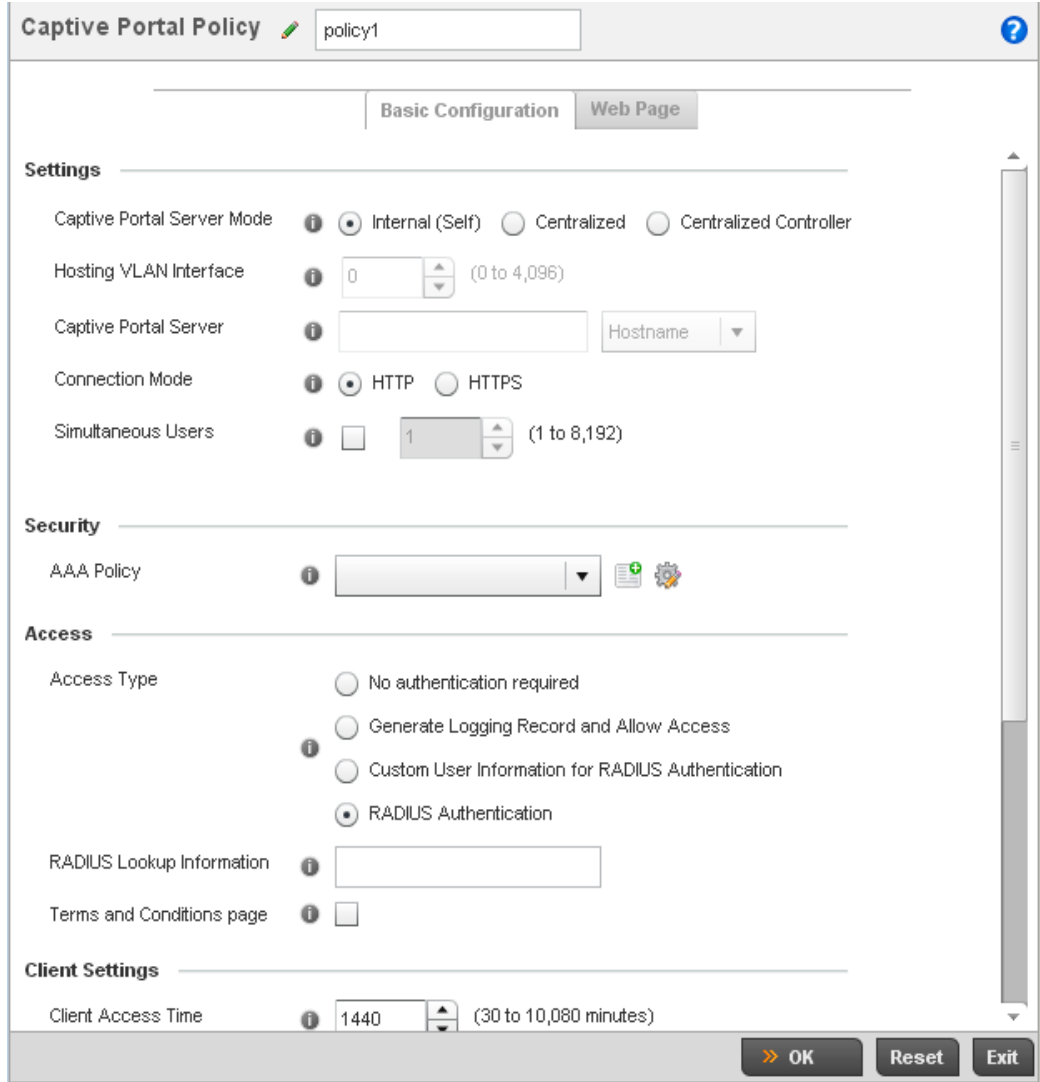


FIGURE 319 Captive Portal Policy screen

3. Refer to the following captive portal policy parameters to determine whether a new policy requires creation, or an existing policy requires edit or deletion:

- Captive Portal Policy** Displays the name assigned to the captive portal guest access policy when it was initially created. A policy name cannot be modified as part of the edit process.
- Captive Portal Server** Lists the IP address (non DNS hostname) of the external (fixed) server validating guest user permissions for the listed captive portal policy. This item remains empty if the controller is hosting the captive portal.
- Captive Portal Server Mode** Lists each policy's hosting mode as either *Internal (Self)* or *External (Fixed)*. If the mode is *Internal (Self)*, the controller is maintaining the captive portal internally, while *External (Fixed)* means the captive portal is being supported on an external server.

- Connection Mode** Lists each policy's connection mode as either HTTP or HTTPS. Both HTTP and HTTPS use the same *Uniform Resource Identifier* (URI), so controller and client resources can be identified. However, Brocade Mobility recommends the use of HTTPS, as it affords transmissions some measure of data protection HTTP cannot provide.
- Simultaneous Users** Displays then number of users permitted at one time for each listed policy. A controller managed captive portal can support between 0-8192 users simultaneously.
- Web Page Source** Displays whether the controller's captive portal HTML pages are maintained *Internally* (on the controller), *Externally* (on an external system you define) or are *Advanced* pages maintained and customized by the network administrator. Internal is the default setting.
- AAA Policy** Lists each controller AAA policy used to authorize client guest access requests. The controller's security provisions provide a way to configure advanced AAA policies that can be applied to captive portal policies. When a captive portal policy is created or modified, a AAA policy must be defined and applied to effectively authorize, authenticate and account user requests.

4. Select **Add** to create a new captive portal policy, **Edit** to modify an existing policy or **Delete** to remove an existing captive portal policy.
5. A **Basic Configuration** screen displays by default. Define the policy's security, access and whitelist basic configuration before actual HTML pages can be defined for guest user access requests.

FIGURE 320 Captive Portal Policy Basic Configuration screen

6. Define the following for the captive portal policy:

Captive Portal Policy	If creating a new policy, assign a name representative of its access permissions, location or intended wireless client user base. If editing an existing captive portal policy, the policy name cannot be modified. The name cannot exceed 32 characters.
Captive Portal Server Mode	Set the mode as either <i>Internal (Self)</i> or <i>External (Fixed)</i> . Select the Internal (Self) radio button for the controller to maintain the captive portal configuration (Web pages) internally. Select the External (Fixed) radio button if the captive portal is supported on an external server. The default value is Internal (Self).
Captive Portal Server	Set a numeric IP address (or DNS hostname) for the server validating guest user permissions for the captive portal policy. This option is only available if hosting the captive portal on an External (Fixed) server resource.
Connection Mode	Select either HTTP or HTTPS to define the connection medium used by the controller and Web serve. Brocade Mobility recommends the use of HTTPS, as it affords the controller some additional data protection HTTP cannot provide. The default value however is HTTP.
Simultaneous Users	Select the checkbox and use the spinner control to set between 1-8192 users (client MAC addresses) allowed to simultaneously access and use the captive portal.

- Use the **AAA Policy** drop-down menu to select the controller *Authentication*, *Authorization* and *Accounting* (AAA) policy used to validate user credentials and provide captive portal guest access to the controller managed network.
- If no AAA policies exist, one must be created by selecting the **Create** icon, or an existing AAA policy can be selected and modified by selecting it from the drop-down menu and selecting the **Edit** icon. For information on creating a AAA policy that can be applied to a captive portal configuration, see *AAA Policy on page 6-289*.
- Set the following **Access** parameters to define hotspot access, RADIUS lookup information and whether the hotspot's login pages contain agreement terms that must be accepted before access is granted to controller resources:

Access Type	Select the radio button for the authentication scheme applied to wireless clients using the captive portal for guest access to the controller managed network. Options include: <i>No authentication required</i> - Clients can freely access the captive portal Web pages without authentication. <i>Generate Logging Record and Allow Access</i> - Access is provided without authentication, but a record of the accessing client is logged. <i>Custom User Information for RADIUS Authentication</i> - When selected, accessing clients are required to provide a 1-32 character lookup data string used to authenticate client access. <i>RADIUS Authentication</i> - An accessing client's user credentials require authentication before access to the captive portal is granted. This is the default setting.
RADIUS Lookup Information	When Custom User Information for RADIUS Authentication is selected as the access type, provide a 1-32 character lookup information string used as a customized authentication mechanism.
Terms and Conditions page	Select this option to include terms that must be adhered to for captive portal access. These terms are included in the Agreement page when <i>No authentication required</i> is selected as the access type, otherwise the terms appear in the Login page. The default setting is disabled.

10. Set the following **Client Settings** to define the duration clients are allowed captive portal access and when they're timed out due to inactivity:

Client Access Time Use the spinner control to define the duration wireless clients are allowed access to the controller managed network using the captive portal policy. Set an interval between 30 - 10,800 minutes. The default interval is 1,440 minutes.

Inactivity Timeout Use the drop-down menu to specify an interval in either *Minutes* (5 - 30) or *Seconds* (300 - 1,800) that, when exceeded, times out clients that have not transmitted a packet within the captive portal.

11. Use the **DNS White List** parameter to create a set of allowed destination IP addresses. These allowed DNS destination IP addresses are called a *Whitelist*.

To effectively host hotspot pages on an external Web server, the IP address of the destination Web server(s) should be in the Whitelist.

12. Refer to the drop-down menu of existing DNS White List entries to select a policy to be applied to this captive portal policy. If no DNS Whitelist entries exist, select the **Create** or **Edit** icons and follow the sub-steps below:
- If creating a new Whitelist, assign it a name up to 32 characters. Use the **+ Add** button to populate the Whitelist with Host and IP Index values.

DNS Entry	Match Suffix
125.89.14.81	✓

FIGURE 321 Captive Portal Whitelist screen

- Provide a numerical IP address or Hostname within the **DNS Entry** parameter for each destination IP address or host included in the Whitelist.
- Use the **Match Suffix** parameter to match any hostname or domain name as a suffix. The default setting is disabled.
- If necessary, select the radio button of an existing Whitelist entry and select the **- Delete** icon to remove the entry from the Whitelist.

13. Set the following **Accounting** parameters to define how accounting is conducted for clients entering and exiting the captive portal. Accounting is the method of collecting and sending security server information for billing, auditing and reporting user data; such as captive portal start and stop times, executed commands (such as PPP), number of packets and number of bytes. Accounting enables wireless network administrators to track captive portal services users are consuming.

Enable RADIUS Accounting

Select this option to use an external RADIUS resource for AAA accounting. When selected, a AAA Policy field displays. This setting is disabled by default.

Enable Syslog Accounting

Select this option to log information about the use of remote access services by users using an external syslog resource. This information is of great assistance in partitioning local versus remote users. Remote user information can be archived to a location outside of the controller for periodic network and user administration. This feature is disabled by default.

Syslog Host

Use the drop-down menu to determine whether an IP address or a host name is used as a syslog host. The IP address or host name of an external server resource is required to route captive portal syslog events to that destination for use as an external controller resource.

Syslog Port

Define the numerical syslog port the controller uses to route traffic with the external syslog server. The default port is 514.

14. Select **OK** to save the changes made within the Basic Configuration screen. Selecting **Reset** reverts the settings back to the last saved configuration.
15. Select the **Web Page** tab to create controller maintained HTML pages requesting wireless clients use to login and navigate within a controller managed hotspot.
16. The **Login** page displays by default.

FIGURE 322 Captive Portal Policy Internal Web Page screen

The *Login* screen prompts the user for a username and password to access the hotspot and proceed to either the Terms and Conditions page (if used) or the Welcome page. The *Terms and Conditions* page provides conditions that must be agreed to before guest access is allowed for the captive portal policy. The *Welcome* page asserts a user has logged in successfully and can access the controller managed hotspot. The *Fail* page asserts the hotspot authentication attempt has failed, and the user is not allowed to access the Internet (using this captive portal) and must provide the correct login information again to access the Internet.

17. Provide the following if creating pages maintained internally on the controller (when the **Internal** radio button is selected as the Web Page Source). The Internal (internally hosted) captive portal is the controller's default setting.
18. The **Login** page displays by default. Configure the following for the **Login**, **Terms and Conditions**, **Welcome** and **Fail** screens:.

Title Text	Set the title text displayed on the Login, Agreement, Welcome and Fail pages when wireless clients access each page. The text should be in the form of a page title describing the respective function of each page and should be unique to each login, agreement, welcome and fail function.
Header Text	Provide header text unique to the function of each page.
Login Message	Specify a message containing unique instructions or information for the users who access the Login, Agreement, Welcome or Fail pages. In the case of the Agreement page, the message can be the conditions requiring agreement before guest access is permitted.

Footer Text	Provide a footer message displayed on the bottom of each page. The footer text should be any concluding message unique to each page before accessing the next page in the succession of hotspot Web pages.
Main Logo URL	The Main Logo URL is the URL for the main logo image displayed on the Login, Agreement, Welcome and Fail pages. Use the Browse button to navigate to the location of the target file.
Small Logo URL	The Small Logo URL is the URL for a small logo image displayed on the Login, Agreement, Welcome and Fail pages. Use the Browse button to navigate to the location of the target file.

19. Select **OK** to save the changes made within the Internal Pages screen. Selecting **Reset** reverts the settings back to the last saved configuration.
20. Select **Advanced** to use a custom directory of Web pages copied to and from the controller for captive portal support.

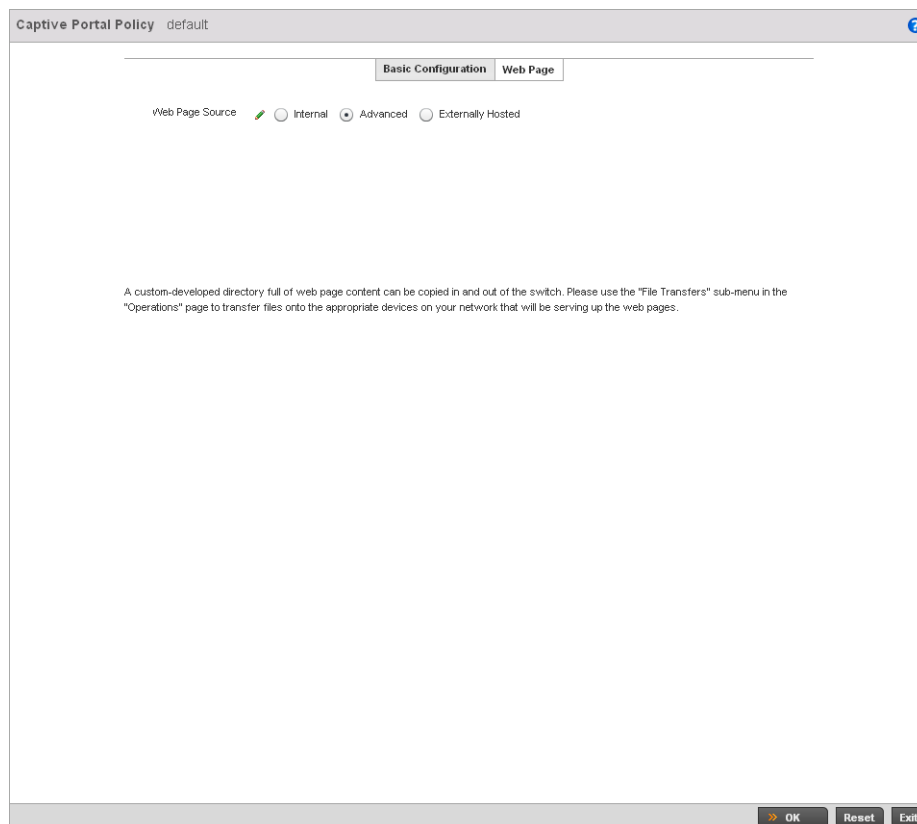


FIGURE 323 Captive Portal Policy Advanced Web Page screen

21. Set the following external URL destinations for the captive portal policy's hotspot pages.

URL	Define the complete URL for the location of the custom captive portal pages.
Advanced	<p>Select the Advanced link to display additional parameters for accessing the remote server used to support the advanced captive portal. The following parameters are required:</p> <p><i>Protocol</i> - Select the file transfer method used between the controller and the resource maintaining the custom captive portal files.</p> <p><i>Port</i> - Use the spinner control to set the port used on the external Server maintaining the custom captive portal files.</p> <p><i>Host</i> - Set the IP address or hostname of the destination server supporting the captive portal's advanced files set. Use the drop-down menu to specify whether an IP address or hostname is used.</p> <p><i>Path</i> - Provide a complete and accurate path to the location where the captive portal file set resides on the external server resource.</p>
Export	Select the Export button to upload target captive portal files to the designated external resource. The exported files display within the File/s table.
Import	Select the Import button to download target captive portal files from the designated external resource to the controller. The imported files display within the File/s table.

22. If hosting the captive portal on a system external to the controller, select the **Externally Hosted** radio button.

Captive Portal Policy default

Basic Configuration Web Page

Web Page Source Internal Advanced Externally Hosted

Login URL

Agreement URL

Welcome URL

Fail URL

A set of pre-existing web pages outside of the switch are specified by the provided URLs.
Three separate URLs point to external web pages for: Logging the user in, Welcoming the user after logging in successfully and Informing the user of a failed login attempt.

OK Reset Exit

FIGURE 324 Captive Portal Policy Externally Hosted Web Page screen

Login URL	Define the complete URL for the location of the Login page. The Login screen prompts the user for a username and password to access the Terms and Conditions or Welcome page.
Agreement URL	Define the complete URL for the location of the Terms and Conditions page. The Terms and Conditions page provides conditions that must be agreed to before wireless client access is provided.
Welcome URL	Define the complete URL for the location of the Welcome page. The Welcome page asserts the user has logged in successfully and can access controller resources via the captive portal.
Fail URL	Define the complete URL for the location of the Fail page. The Fail page asserts authentication attempt has failed, and the client cannot access the captive portal and the client needs to provide correct login information to regain access.

23. Select **OK** when completed to update the captive portal's advanced configuration. Select **Reset** to revert the screen back to its last saved configuration.

Creating DNS Whitelists

Before defining a captive portal configuration, refer to the following deployment guidelines to ensure the configuration is optimally effective:

1. Select **Configuration > Services**.

The upper, left-hand, side of the user interface displays a **Services** menu pane where Captive Portal, DHCP and RADIUS configuration options can be selected.

2. Select **Captive Portals**.

The Captive Portal screen displays the configurations of existing policies. New policies can be created, existing policies can be modified or existing policies deleted.

3. Select **DNS Whitelist**




FIGURE 325 Captive Portal DNS Whitelist screen

4. Review the names of existing DNS Whitelists and click **Add** to create a new whitelist entry or select an existing whitelist and click **Edit** to modify it.

5. Use the **DNS Whitelist** parameter to create a set of allowed destination IP addresses. These allowed DNS destination IP addresses are called a *Whitelist*.

To effectively host hotspot pages on an external Web server, the IP address of the destination Web server(s) should be in the Whitelist.

6. Refer to the drop-down menu of existing DNS White List entries to select a policy to be applied to this captive portal policy. If no DNS Whitelist entries exist, select the **Create** or **Edit** icons and follow the sub-steps below:

- a. If creating a new Whitelist, assign it a name up to 32 characters. Use the **+ Add** button to populate the Whitelist with Host and IP Index values.

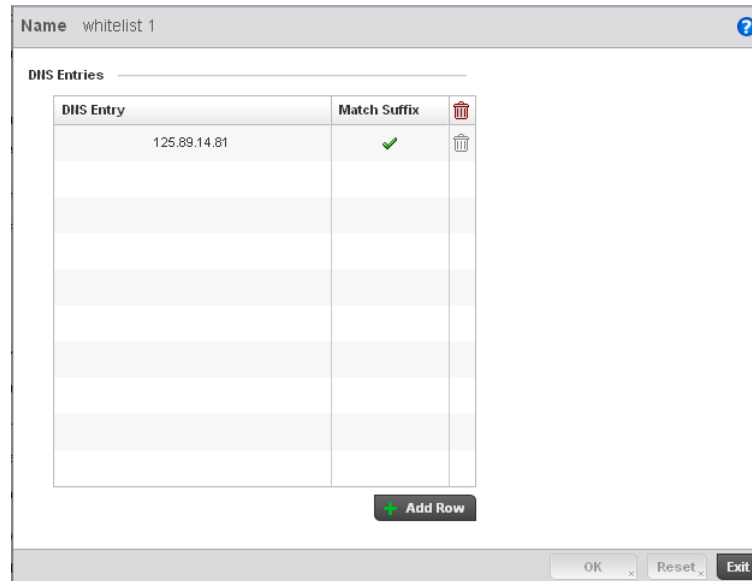


FIGURE 326 Captive Portal Whitelist screen

- b. Provide a numerical IP address or Hostname within the **DNS Entry** parameter for each destination IP address or host included in the Whitelist.
- c. Use the **Match Suffix** parameter to match any hostname or domain name as a suffix. The default setting is disabled.
- d. If necessary, select the radio button of an existing Whitelist entry and select the - **Delete** icon to remove the entry from the Whitelist.

Captive Portal Deployment Considerations

Before defining a captive portal configuration, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- The architecture should consider the number of wireless clients allowed on the guest network and the services provided. Each topology has benefits and disadvantages which should be taken into consideration to meet each deployment's requirements.
- Hotspot authentication uses secure HTTPS to protect user credentials, but doesn't typically provide encryption for user data once they have been authenticated. For private access applications, Brocade Mobility recommends WPA2 (with a strong passphrase) be enabled to provide strong encryption.
- Brocade Mobility recommends guest user traffic be assigned a dedicated VLAN, separate from other internal networks.
- Controller guest access configurations should include firewall policies to ensure logical separation is provided between guest and internal networks so internal networks and hosts are not reachable from controller managed guest devices.
- Guest access services should be defined in a manner whereby end-user traffic doesn't cause network congestion.

- Brocade Mobility recommends a valid certificate be issued and installed on all devices providing Hotspot access to a controller managed WLAN and wireless network. The certificate should be issued from a public certificate authority ensuring guests can access the Hotspot without browser errors.

Setting the Controller's DHCP Configuration

Dynamic Host Configuration Protocol (DHCP) allows hosts on an IP network to request and be assigned IP addresses and discover information about the controller managed network where they reside. Each subnet can be configured with its own address pool. Whenever a DHCP client requests an IP address, the DHCP server assigns an IP address from that subnet's address pool. When the controller's onboard DHCP server allocates an address for a DHCP client, the client is assigned a lease, which expires after an pre-determined interval. Before a lease expires, wireless clients (to which leases are assigned) are expected to renew them to continue to use the addresses. Once the lease expires, the client is no longer permitted to use the leased IP address. The controller's DHCP server ensures all IP addresses are unique, and no IP address is assigned to a second client while the first client's assignment is valid (its lease has not yet expired). Therefore, IP address management is conducted by the controller's DHCP server, not by an administrator.

The controller's internal DHCP server groups wireless clients based on defined user-class options. Clients with a defined set of user class values are segregated by class. A DHCP server can associate multiple classes to each pool. Each class in a pool is assigned an exclusive range of IP addresses. DHCP clients are compared against classes. If the client matches one of the classes assigned to the pool, it receives an IP address from the range assigned to the class. If the client doesn't match any of the classes in the pool, it receives an IP address from a default pool range (if defined). Multiple IP addresses for a single VLAN allow the configuration of multiple IP addresses, each belonging to different subnet. Class configuration allows a DHCP client to obtain an address from the first pool to which the class is assigned.

NOTE

DHCP server updates are only implemented when the controller is restarted.

To access and review the controller's internal DHCP server configuration:

1. Select **Configuration > Services > DHCP Server Policy**.
2. The **DHCP Server** screen displays. Clients with a defined set of user class values are segregated by class. A DHCP server can associate multiple classes to each pool. Each class in a pool is assigned an exclusive range of IP addresses. DHCP clients are then compared against classes.

- Review the following DHCP pool configurations to determine if an existing pool can be used as is, a new one requires creation or edit, or a pool requires deletion:

DHCP Pool	Displays the name assigned to the network pool when created. The DHCP pool name represents the group of IP addresses used to assign to DHCP clients upon request. The name assigned cannot be modified as part of the edit process. However, if the network pool configuration is obsolete it can be deleted.
Subnet	Displays the network address and mask used by clients requesting DHCP resources.
Domain Name	Displays the domain name defined used by the controller with this network pool. The controller uses <i>Domain Name Service</i> (DNS) to convert human readable host names into IP addresses. Host names are not case sensitive and can contain alphabetic or numeric letters or a hyphen. A <i>fully qualified domain name</i> (FQDN) consists of a host name plus a domain name. For example, <i>computername.domain.com</i> .
Boot File	Boot files (<i>Boot Protocol</i>) are used to boot remote systems over the controller managed network. BOOTP messages are encapsulated inside UDP messages so requests and replies can be forwarded by the controller. Each DHCP network pool can use a different file as needed.
Lease Time	If a lease time has been defined for a listed network pool, it displays in an interval between 1 - 31,622,399 seconds. DHCP leases provide addresses for defined times to various clients. If a client does not use the leased address for the defined time, that IP address can be re-assigned to another DHCP supported client.

- Select **Add** to create a new DHCP pool, **Edit** to modify an existing pool's properties or **Delete** to remove a pool from amongst those available.

The screenshot shows the DHCP Pool configuration interface for a pool named 'pool1'. The 'Basic Settings' tab is active. The 'General' section includes a Subnet dropdown set to '255.255.255.0 / 24', an empty Domain Name field, and a DNS Servers table with six rows of '0.0.0.0' and 'Clear' buttons. The 'Lease Time' is checked and set to '86400'. The 'Default Routers' field is empty, and there is another table with six rows of '0.0.0.0' and 'Clear' buttons. Below these are two empty tables for 'IP Address Ranges' and 'Excluded IP Address Range', each with columns for 'IP Start', 'IP End', and 'Class Policy', and an 'Add Row' button. The bottom of the screen has 'OK', 'Reset', and 'Exit' buttons.

FIGURE 329 DHCP Pools screen - Basic Settings tab

If adding or editing a DHCP pool, the DHCP Pool screen displays the **Basic Settings** tab by default. Define the required parameters for the Basic Settings, Static Bindings and Advanced tabs to complete the creation of the DHCP pool.

4. Set the following **General** parameters from within the **Basic Settings** tab:

- DHCP Pool** If adding a new pool, a name is required. The pool is the range of IP addresses defined for DHCP assignment or lease. The name assigned cannot be modified as part of the edit process. However, if the network pool configuration is obsolete it can be deleted. The name cannot exceed 32 characters.
- Subnet** Define the IP address and Subnet Mask used for DHCP discovery and requests between the controller's DHCP Server and DHCP clients. The IP address and subnet mask are required to match the addresses of the layer 3 interface for the addresses to be supported through that interface.
- Domain Name** Provide the domain name used by the controller with this pool. Domain names are not case sensitive and can contain alphabetic or numeric letters or a hyphen. A *fully qualified domain name* (FQDN) consists of a host name plus a domain name. For example, *computername.domain.com*.

DNS Servers	Define one or a group of <i>Domain Name Servers</i> (DNS) to translate domain names to IP addresses. Select clear to remove any single IP address as needed. Up to 8 IP addresses can be supported.
Lease Time	DHCP leases provide addresses for defined times to various clients. If a client does not use the leased address for the defined time, that IP address can be re-assigned to another DHCP supported client. Select this option to assign a lease in either <i>Seconds</i> (1 - 31, 622, 399), <i>Minutes</i> (1 - 527,040), <i>Hours</i> (1 - 8,784) or <i>Days</i> (1 - 366). The default setting is enabled, with a lease time of 1 day.
Default Routers	After a DHCP client has booted, the client begins sending packets to its default router. Set the IP address of one or a group of routers the controller uses to map host names into IP addresses available to DHCP supported clients. Up to 8 default router IP addresses are supported.

5. Use the **IP Address Ranges** field define the range of included (starting and ending IP addresses) addresses for this particular pool.
 - a. Select the **+ Add Row** button at the bottom of the IP addresses field to add a new range. At any time you can select the radio button of an existing IP address range and select the **Delete** icon to remove it from the list of those available.
 - b. Enter a viable range of IP addresses in the **IP Start** and **IP End** columns. This is the range of addresses available for assignment to DHCP supported wireless clients within the controller managed network.
 - c. Select the **Create** icon or **Edit** icon within the **Class Policy** column to display the **DHCP Server Policy** screen if a class policy is not available from the drop-down menu.
 - d. Refer to the **Excluded IP Address Range** field and select the **+Add Row** button. Add ranges of IP address to exclude from lease to requesting DHCP clients. Having ranges of unavailable addresses is a good practice to ensure IP address resources are in reserve. Select the **Delete** icon as needed to remove an excluded address range.
 - e. Select **OK** to save the updates to the DHCP Pool Basic Settings tab. Select **Reset** to revert to the last saved configuration.

6. Select the **Static Bindings** tab from within the DHCP Pools screen.

A binding is a collection of configuration parameters, including an IP address, associated with, or *bound to*, a DHCP client. Bindings are managed by DHCP servers. DHCP bindings automatically map a device MAC address to an IP address using a pool of DHCP supplied addresses. Static bindings provide the assignment of IP addresses without creating numerous host pools with manual bindings. Static host bindings use a text file the DHCP server reads. It eliminates the need for a lengthy configuration file and reduces the space required on the controller to maintain address pools.

Client Identifier Type hardware-address Value 010203040506

General

IP Address * 192 . 8 . 7 . 4

Domain Name

Boot File

BOOTP Next Server

Client Name

Enable Unicast

Static Routes Installed on Clients

Destination	Gateway

Add Row

DHCP Option Values

Global DHCP Option Name	Value

Add Row

NetBIOS

NetBIOS Node Type Undefined

NetBIOS Servers

IP Address
0 . 0 . 0 . 0 Clear
0 . 0 . 0 . 0 Clear
0 . 0 . 0 . 0 Clear
0 . 0 . 0 . 0 Clear
0 . 0 . 0 . 0 Clear

Network

DNS Servers

IP Address
0 . 0 . 0 . 0 Clear
0 . 0 . 0 . 0 Clear
0 . 0 . 0 . 0 Clear
0 . 0 . 0 . 0 Clear
0 . 0 . 0 . 0 Clear

Default Routers

IP Address
0 . 0 . 0 . 0 Clear
0 . 0 . 0 . 0 Clear
0 . 0 . 0 . 0 Clear
0 . 0 . 0 . 0 Clear
0 . 0 . 0 . 0 Clear

OK

FIGURE 331 Static Bindings Add screen

9. Define the following **General** parameters to complete the creation of the static binding configuration:

IP Address	Set the IP address of the client using this host pool.
Domain Name	Provide a domain name of the current interface. Domain names aren't case sensitive and can contain alphabetic or numeric letters or a hyphen. A <i>fully qualified domain name</i> (FQDN) consists of a host name plus a domain name. For example, <i>computername.domain.com</i>
Boot File	Enter the name of the boot file used with this pool. Boot files (Boot Protocol) can be used to boot remote systems over the controller managed network. BOOTP messages are encapsulated inside UDP messages so requests and replies can be forwarded by the controller. Each DHCP network pool can use a different file as needed
BOOTP Next Server	Provide the numerical IP address of the server providing BOOTP resources.
Client Name	Provide the name of the client requesting DHCP Server support.
Enable Unicast	Unicast packets are sent from one location to another location (there's just one sender, and one receiver). Select this option to allow the controller to forward unicast messages to just a single device within this network pool.

10. Define the following **NetBIOS** parameters to complete the creation of the static binding configuration:

NetBIOS Node Type	Set the NetBios Node Type used with this particular pool. The node can have one of the following types: <i>Broadcast</i> - Uses broadcasting to query nodes on the network for the owner of a NetBIOS name. <i>Peer-to-Peer</i> - Uses directed calls to communicate with a known NetBIOS name server (such as a WINS server), for the IP address of a NetBIOS machine. <i>Mixed</i> - A mixed node using broadcasted queries to find a node, and failing that, queries a known p-node name server for the address. <i>Hybrid</i> - A combination of two or more nodes. <i>Undefined</i> - No node type is applied.
NetBIOS Servers	Specify a numerical IP address of a single or group of NetBIOS WINS servers available to DHCP supported wireless clients. A maximum of 8 server IP addresses can be assigned.

11. Refer to the **Static Routes Installed on Clients** field to set **Destination IP** and **Gateway** addresses enabling assignment of static IP addresses without creating numerous host pools with manual bindings. This eliminates the need for a long configuration file and reduces the space required in NVRAM to maintain address pools. Select the **+ Add Row** button to add individual destinations. Select the **Delete** icon to remove it from the list of those available.
12. Refer to the **DHCP Option Values** table to set Global DHCP options. A set of global DHCP options applies to all clients, whereas a set of subnet options applies only to the clients on a specified subnet. If you configure the same option in more than one set of options, the precedence of the option type decides which the DHCP server supports a client.
- Select the **+ Add Row** button to add individual options. Assign each a **Global DHCP Option Name** to help differentiate it from others with similar configurations. At any time you can select the radio button of an existing option and select the **- Delete** button to remove it from the list of those available.
 - Assign a **Value** to each option with codes in the range 1 through 254. A vendor-specific option definition only applies to the vendor class for which it is defined.
13. Within the **Network** section, define one or group of **DNS Servers** to translate domain names to IP addresses. Up to 8 IP addresses can be provided.
14. Select **OK** when completed to update the static bindings configuration. Select **Reset** to revert the screen back to its last saved configuration.
15. Select the **Advanced** tab to define additional NetBIOS and Dynamic DNS parameters.

FIGURE 332 DHCP Pools screen - Advanced tab

16. The addition or edit of the network pool's advanced settings requires the following **General** parameters be set:

- Boot File** Enter the name of the boot file used with this pool. Boot files (Boot Protocol) can be used to boot remote systems over the controller managed network. BOOTP messages are encapsulated inside UDP messages so requests and replies can be forwarded by the controller. Each pool can use a different file as needed.
- BOOTP Next Server** Provide the numerical IP address of the server providing BOOTP resources.
- Enable Unicast** Unicast packets are sent from one location to another location (there's just one sender, and one receiver). Select this option to allow the controller to forward unicast messages to just a single device within the network pool.

17. Set the following **NetBIOS** parameters for the network pool:

- NetBIOS Node Type** Set the NetBIOS Node Type used with this pool. The following types are available:
Broadcast - Uses broadcasting to query nodes on the network for the owner of a NetBIOS name.
Peer-to-Peer - Uses directed calls to communicate with a known NetBIOS name server, such as a WINS server, for the IP address of a NetBIOS machine.
Mixed - Is a mixed node using broadcasted queries to find a node, and failing that, queries a known p-node name server for the address.
Hybrid - Is a combination of two or more nodes.
Undefined - No NetBIOS Node Type is used.
- NetBIOS Servers** Specify a numerical IP address of a single or group of NetBIOS WINS servers available to DHCP supported wireless clients.

18. Define the following set of **Dynamic DNS (Not Applicable for Static Bindings)** parameters used with the network pool configuration. DDNS enables the controller to notify a DNS server to change, in real time (ad-hoc) the active DNS configuration of its configured hostnames, addresses or other information stored in DNS.

DDNS Domain Name	Enter a domain name for DDNS updates representing the forward zone in the DNS server. For example, <i>test.net</i> .
DDNS TTL	Select this option to set a TTL (Time to Live) to specify the validity of DDNS records. The maximum value configurable is 864000 seconds.
DDNS Multi User Class	Select the check box to associate the user class option names with a multiple user class. This allows the user class to transmit multiple option values to DHCP servers supporting multiple user class options.
Update DNS	Set if DNS is updated from a client or a server. Select either <i>Do Not Update</i> , <i>Update from Server</i> or <i>Update from Client</i> . The default setting is <i>Do Not Update</i> , implying that no DNS updates occur at all.
DDNS Server	Specify a numerical IP address of one or two DDNS servers.

19. Refer to the **DHCP Option Values** table to set global DHCP options applicable to all clients, whereas a set of subnet options applies to just the clients on a specified subnet.
- Select the **+ Add Row** button to add individual options. Assign each a **Global DHCP Option Name** to help differentiate it from others with similar configurations. At any time you can select the radio button of an existing option and select the **Delete** icon to remove it from the list of those available.
 - Assign a **Value** to each option with codes in the range 1 through 254. A vendor-specific option definition only applies to the vendor class for which it's defined.
20. Click the **+ Add Row** button and enter a **Destination** and **Gateway IP Address** to add **Static Routes Installed on Clients**.
21. Select **OK** to save the updates to the DHCP pool's Advanced settings. Select **Reset** to revert the screen back to its last saved configuration.

Defining DHCP Server Global Settings

Setting a DHCP server global configuration entails defining whether BOOTP requests are ignored and setting DHCP global server options.

To define DHCP server global settings:

- Select **DHCP > Server Policy** from within Services menu pane.
- Select the **Global Settings** tab.

The screenshot shows the 'Global Settings' tab of the DHCP Server Policy configuration. The window title is 'DHCP Server Policy test dhcp'. There are three tabs: 'DHCP Pool', 'Global Settings', and 'Class Policy'. The 'Global Settings' tab is active.

Configuration

Ignore BOOTP Requests:

Ping Timeout: seconds (1 to 10)

Global DHCP Server Options

Name	Type	Code	

Buttons: + Add Row, OK, Reset, Exit

FIGURE 333 DHCP Server Policy screen - Global Settings tab

3. Set the following parameters within the **Configuration** field:

Ignore BOOTP Requests Select the checkbox to ignore BOOTP requests. BOOTP (boot protocol) requests boot remote systems within the controller managed network. BOOTP messages are encapsulated inside UDP messages and are forwarded by the controller. This feature is disabled by default, so unless selected, BOOTP requests are forwarded.

Ping Timeout Set an interval (from 1 -10 seconds) for the DHCP server ping timeout. The controller uses the timeout to intermittently ping and discover whether a client requested IP address is already used.

4. Refer to the **Global DHCP Server Options** field.

- Use the **+ Add Row** button at the bottom of the field to add a new global DHCP server option. At any time you can select the radio button of an existing global DHCP server option and select the **Delete** icon to remove it from the list of those available.
- Use the **Type** drop-down menu to specify whether the DHCP option is being defined as a numerical IP address or ASCII string or Hex string. Highlight an entry from within the Global Options screen and click the Remove button to delete the name and value.

5. Select **OK** to save the updates to the DHCP server global settings. Select **Reset** to revert the screen back to its last saved configuration.

DHCP Class Policy Configuration

The controller's DHCP server assigns IP addresses to DHCP enabled wireless clients based on user class option names. Clients with a defined set of user class option names are identified by their user class name. The DHCP server can assign IP addresses from as many IP address ranges as defined by the administrator. The DHCP user class associates a particular range of IP addresses to a device in such a way that all devices of that type are assigned IP addresses from the defined range.

Refer to the **DHCP Class Policy** screen to review existing DHCP class names and their current multiple user class designations. Multiple user class options enable a user class to transmit multiple option values to DHCP servers supporting multiple user class options. Either add a new class policy, edit the configuration of an existing policy or permanently delete a policy as required.

To review DHCP class policies:

1. Select **Configuration > Services**.
2. Select the **Class Policy** tab.

The screenshot shows the DHCP Server Policy configuration interface. At the top, there are tabs for 'DHCP Pool', 'Global Settings', and 'Class Policy'. Below the tabs is a table with the following columns: 'DHCP Class Name' and 'Multiple User Class Support'. The table contains one row with the class name 'user1' and a green checkmark in the 'Multiple User Class Support' column. At the bottom of the table, there is a search bar labeled 'Type to search in tables' and a 'Row Count: 1' indicator. Below the table are buttons for 'Add', 'Edit', 'Delete', and 'Exit'.

DHCP Class Name	Multiple User Class Support
user1	✓

FIGURE 334 DHCP Server Policy screen - Class Policy tab

3. Refer to the following to determine whether a new class policy requires creation, an existing class policy requires edit or an existing policy requires deletion:

DHCP Class Name Displays client names grouped by the class name assigned when the class policy was created.

Multiple User Class A green check mark in this column defines multiple user class support as enabled from the listed DHCP class name. A red "X" defines multiple user class support as disabled for the listed DHCP class name. Multiple user class support can be *enabled/disabled* for existing class names by editing the class name's configuration.

4. Select **Add** to create a new DHCP class policy, **Edit** to update an existing policy or **Delete** to remove an existing policy.

FIGURE 335 DHCP Class Name Add screen

5. If adding a new **DHCP Class Name**, assign a name representative of the device class supported. The DHCP user class name should not exceed 32 characters.
6. Select a row within the **Value** column to enter a 32 character maximum value string.
7. Select the **Multiple User Class** check box to enable multiple option values for the user class. This allows the user class to transmit multiple option values to DHCP servers supporting multiple user class options.
8. Select **OK** to save the updates to this DHCP class policy. Select **Reset** to revert the screen back to its last saved configuration.

DHCP Deployment Considerations

Before defining a controller's internal DHCP server configuration, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- Brocade Mobility DHCP option 189 is required when Brocade Mobility 650 Access Point access points are deployed over a layer 3 network and require layer 3 adoption. DHCP services are not required for Brocade Mobility 650 Access Point access points connected to a VLAN that's local to the controller.

- DHCP's lack of an authentication mechanism means a DHCP server cannot check if a client or user is authorized to use a given user class. This introduces a vulnerability when using user class options. For example, if a user class is used to assign a special parameter (for example, a database server), there is no way to authenticate a client and it's impossible to check if a client is authorized to use this parameter.
- Ensure traffic can pass on UDP ports 67 and 68 between the controller and the clients receiving DHCP information to ensure optimum DHCP support for DHCP enabled wireless clients.

Setting the Controller's RADIUS Configuration

Remote Authentication Dial-In User Service (RADIUS) is a client/server protocol and software enabling remote access servers to communicate with the switch to authenticate users and authorize their access to the controller managed network. RADIUS is a distributed client/server system that secures networks against unauthorized access. RADIUS clients send authentication requests to the controller's RADIUS server containing user authentication and network service access information.

RADIUS enables centralized management of controller authentication data (usernames and passwords). When an MU attempts to associate to the RADIUS supported controller, the controller sends the authentication request to the RADIUS server. The authentication and encryption of communications between the controller and server takes place through the use of a shared secret password (not transmitted over the network).

The controller's local RADIUS server stores the user database locally, and can optionally use a remote user database. It ensures higher accounting performance. It allows the configuration of multiple users, and assign policies for the group authorization.

The controller allows the enforcement of user-based policies. User policies include dynamic VLAN assignment and access based on time of day. The controller uses a default trustpoint. A certificate is required for EAP TTLS, PEAP and TLS RADIUS authentication (configured with the RADIUS service).

Dynamic VLAN assignment is achieved based on the RADIUS server response. A user who associates to WLAN1 (mapped to VLAN1) can be assigned a different VLAN after authentication with the RADIUS server. This dynamic VLAN assignment overrides the WLAN's VLAN ID to which the User associates.

To view RADIUS configurations:

Select **Configuration** tab from the main menu.

Select the **Services** tab from the **Configuration** menu.

The upper, left-hand side pane of the User interface displays the **RADIUS** option. The **RADIUS Group** screen displays (by default).

For information on creating the groups, user pools and server policies needed to validate user credentials against a server policy configuration, refer to the following:

- [Creating RADIUS Groups](#)
- [Defining User Pools](#)
- [Configuring RADIUS Server Policies](#)
- [RADIUS Deployment Considerations](#)

Creating RADIUS Groups

The controller's RADIUS server allows the configuration of user groups with common user policies. User group names and associated users are stored in the controller's local database. The user ID in the received access request is mapped to the associated wireless group for authentication. The configuration of groups allows the enforcement of the following policies that can control the access of the user:

- Assign a VLAN to the user upon successful authentication
- Define a start and end of time in (HH:MM) when the user is allowed to authenticate
- Define the list of SSIDs to which a user belonging to this group is allowed to associate
- Define the days of the week the user is allowed to login
- Rate limit traffic

To access RADIUS Groups menu:

1. Select **Configuration** tab from the main menu.
2. Select the **Services** tab from the **Configuration** menu.
3. Select **RADIUS > Groups** from the **Configuration > Services** menu.
4. The browser displays a list of the existing groups.

RADIUS Group	Guest User Group	Management Group	Role	VLAN	Time Start	Time Stop
myGuestGroup	✓	✗		1	12:00 am	11:59 pm

Type to search in tables

Row Count: 1

Add Edit Delete

FIGURE 336 RADIUS Group screen

5. Select a group from the **Group Browser** to view the following read-only information for existing groups:

RADIUS Group	Displays the group name or identifier assigned to each listed group when it was created. The name cannot exceed 32 characters or be modified as part of the group edit process.
Guest User Group	Specifies whether a user group only has guest access and temporary permissions to the controller's local RADIUS server. The terms of the guest access can be set uniquely for each group. A red "X" designates the group as having permanent access to the local RADIUS server. Guest user groups cannot be made management groups with unique access and role permissions.
Management Group	A green checkmark designates this RADIUS user group as a management group. Management groups can be assigned unique access and role permissions.
Role	If a group is listed as a management group, it may also have a unique role assigned. Available roles include: <i>monitor</i> - Read-only access. <i>helpdesk</i> - Helpdesk/support access <i>network-admin</i> - Wired and wireless access <i>security-admin</i> - Grants full read/write access <i>system-admin</i> - System administrator access
VLAN	Displays the VLAN ID used by the group. The VLAN ID is representative of the shared SSID each group member (user) employs to interoperate within the controller managed network (once authenticated by the local RADIUS server).
Time Start	Specifies the time users within each listed group can access the controller's local RADIUS resources.
Time Stop	Specifies the time users within each listed group lose access to the controller's local RADIUS resources.

6. To modify the settings of an existing group, select the group and click the **Edit** button. To delete an obsolete group, select the group and click the **Delete** button.

Creating RADIUS Groups

To create a RADIUS group:

1. Select **Configuration** tab from the main menu.
2. Select the **Services** tab from the **Configuration** menu.
3. Select **RADIUS > Groups** from the **Configuration > Services** menu.
4. Click the **Add** to create a new RADIUS group, **Edit** to modify the configuration of an existing group or **Delete** to permanently remove a selected group.

FIGURE 337 RADIUS Group Policy Add screen

5. Define the following Settings to define the user group configuration:

- RADIUS Group Policy** If creating a new RADIUS group, assign it a name to help differentiate it from others with similar configurations. The name cannot exceed 32 characters or be modified as part of a RADIUS group edit process.
- Guest User Group** Select this option to assign only guest access and temporary permissions to the controller's local RADIUS server. Guest user groups cannot be made management groups with unique access and role permissions.
- VLAN** Select this option to assign a specific VLAN to this RADIUS user group. Ensure Dynamic VLAN assignment (Single VLAN) is enabled for the WLAN in order for the VLAN assignment to work properly. For more information, see *Basic WLAN Configuration on page 6-229*.
- WLAN SSID** Assign a list of SSIDs users within this RADIUS group are allowed to associate to. An SSID cannot exceed 32 characters. Assign WLAN SSIDs representative of the configurations a guest user will need to access. The parameter is not available if this RADIUS group is a management group.
- Rate Limit from Air** Select the checkbox to set the rate limit to controller managed clients within this RADIUS group. Use the spinner to set value from 100-1,000,000 kbps. Setting a value of 0 disables rate limiting

Rate Limit to Air	Select the checkbox to set the rate limit from clients within this RADIUS group. Use the spinner to set value from 100-1,000,000 kbps. Setting a value of 0 disables rate limiting.
Management Group	Select this option to designate this RADIUS group as a management group. This feature is disabled by default. If set as management group, assign a role to the members of the group using the Access drop-down menu allowing varying levels of administrative rights.
Role	If a group is listed as a management group, it may also have a unique role assigned. Available roles include: <i>monitor</i> - Read-only access. <i>helpdesk</i> - Helpdesk/support access <i>network-admin</i> - Wired and wireless access <i>security-admin</i> - Grants full read/write access <i>system-admin</i> - System administrator access

6. Set the **Schedule** to configure access times and dates.

Time Start	Use the spinner control to set the time (in HH:MM format) RADIUS group members are allowed to login and access the controller's RADIUS server resources (for example, 14:45 = 2:45). Select either the AM or PM radio button to set the time as morning or evening.
Time Stop	Use the spinner control to set the time (in HH:MM format) RADIUS group members are denied access to the controller's RADIUS server resources (for example, 15:45 = 3:45). Select either the AM or PM radio button to set the time as morning or evening. If already logged in, the RADIUS group user is deauthenticated from the WLAN.
Days	Select the day(s) of the week RADIUS group members can access controller RADIUS resources.

7. Click the **OK** to save the changes. Select **Reset** to revert to the last saved configuration.

Defining User Pools

A user pool defines policies for individual user access to the controller's internal RADIUS resources. User or pools provide a convenient means of providing user access to the controller's RADIUS resources based on the pool's unique permissions (either temporary or permanent). A pool can contain a single user or group of users.

To configure a RADIUS user pool and unique user IDs:

1. Select **Configuration** from the main menu.
2. Select **Services** tab from the Configuration screen.
3. Select **RADIUS > User Pools** from the **Configuration > Services** menu.




FIGURE 338 RADIUS User Pool screen

4. The **RADIUS User Pool** screen lists the default pool along with any other admin created user pool.
5. Select **Add** to create a new user pool, **Edit** to modify the configuration of an existing pool or **Delete** to remove a selected pool.
6. If creating a new pool, assign it a name up to 32 characters and select **Continue**. The name should be representative of the users comprising the pool and/or the temporary or permanent access privileges assigned.

FIGURE 340 RADIUS User screen

9. Refer the following fields in the **User** screen to create a new user Id with unique access privileges:

User Id	Assign a unique alphanumeric string identifying this user. The Id cannot exceed 64 characters in length.
Password	Provide a password unique to this user Id. The password cannot exceed 32 characters in length. Select the Show checkbox to expose the password's actual character string, leaving the option unselected displays the password as a string of asterisks (*).
Guest User	Select the checkbox to designate this user as a guest with temporary access. The guest user must be assigned unique access times to restrict their access.
Group List	If the user Id has been defined as a guest, use the Group drop-down menu to assign the user a group with temporary access privileges. If the user is defined as a permanent user, select a group from the group list. If there's no groups listed relevant to the user's intended access, select the Create link (or icon for guests) and create a new group configuration suitable for the user Id's membership. For more information, see <i>Creating RADIUS Groups on page 10-493</i> .

10. Select **OK** to save the user Id's group membership configuration. Select **Reset** to revert to the last saved configuration.

Configuring RADIUS Server Policies

A RADIUS server policy is a unique authentication and authorization configuration for receiving user connection requests, authenticating users and returning the configuration information necessary for the RADIUS client to deliver service to the user. The client is the entity with authentication information requiring validation. The controller's RADIUS server has access to a database of authentication information used to validate the client's authentication request.

The controller's RADIUS server ensures the information is correct using authentication schemes like PAP, CHAP or EAP. The user's proof of identification is verified, along with, optionally, other information. A controller's RADIUS server policy can also be configured to refer to an external LDAP resource to verify the user's credentials.

To review RADIUS existing server policies, manage the creation of new policies or manage the modification of existing policies:

1. Select **Configuration** from the main menu.
2. Select **Services** tab from the Configuration screen.
3. Select **RADIUS > Server Policy** from the **Configuration > Services** menu.
4. The **Server Policy Browser** lists existing server policies by either by group or randomly as defined using the drop-down menu. A policy can be selected and modified at any time from the browser.

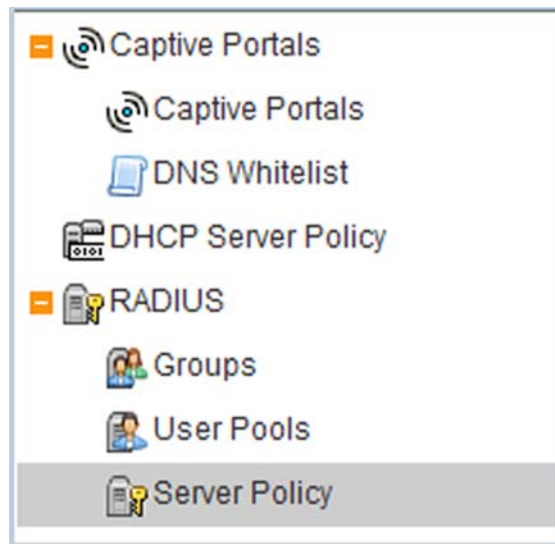


FIGURE 341 Server Policy browser

5. Refer to the RADIUS Server screen to review high-level server policy configuration data.

Local Authentication Type

Lists the controllers local EAP authentication scheme used with this policy. The following EAP authentication types are supported by the controller's onboard RADIUS server:

All - Enables both TTLS and PEAP.

TLS - Uses TLS as the EAP type

TLS and MD5 - The EAP type is TTLS with default authentication using MD5.

TTLS and PAP - The EAP type is TTLS with default authentication using PAP.

TTLS and MSCHAPv2 - The EAP type is TTLS with default authentication using MSCHAPv2.

PEAP and GTC - The EAP type is PEAP with default authentication using GTC.

PEAP and MSCHAPv2 - The EAP type is PEAP with default authentication using MSCHAPv2.

LDAP Authentication Type

Lists the LDAP authentication scheme used with this policy. The following LDAP authentication types are supported with the controller's external LDAP resource:

All - Enables both TTLS and PAP and PEAP and GTC.

TTLS and PAP - The EAP type is TTLS with default authentication using PAP.

PEAP and GTC - The EAP type is PEAP with default authentication using GTC.

CRL Validation

Specifies whether a *Certificate Revocation List* (CRL) check is made. A green checkmark indicates CRL validation is enabled. A red "X" indicates it's disabled. A CRL is a list of revoked certificates issued and subsequently revoked by a *Certification Authority* (CA). Certificates can be revoked for a number of reasons including failure or compromise of a device using a certificate, a compromise of a certificate key pair or errors within an issued certificate. The mechanism used for certificate revocation depends on the CA.

7. Select either **Add** to create a new RADIUS server policy, **Edit** to modify an existing policy or **Delete** to permanently remove a policy.

The screenshot shows the 'RADIUS Server Policy' configuration window for 'policy1'. It has tabs for 'Server Policy', 'Client', 'Proxy', and 'LDAP', with 'Server Policy' selected. The 'Settings' section contains:

- RADIUS User Pools:** A list with checkboxes for 'defaultPool' and 'anotherPool', and a 'Create' button.
- LDAP Server Dead Period:** A text box with '5' and a 'Minutes' dropdown, with '(0 to 10)' in parentheses.
- LDAP Groups:** A dropdown menu showing '<none>', with 'Create' and 'Edit' icons.
- LDAP Group Verification:** A checked checkbox.
- Local Realm:** A text box with 'Create' and 'Edit' icons.

 The 'Authentication' section contains:

- Authentication Data Source:** Radio buttons for 'Local' (selected) and 'LDAP'.
- Local Authentication Type:** A dropdown menu showing 'All'.
- LDAP Authentication Type:** A dropdown menu showing 'All'.
- Enable CRL Validation:** An unchecked checkbox.

 The 'Session Resumption / Fast Reauthentication' section contains:

- Enable Session Resumption:** An unchecked checkbox.
- Cached Entry Lifetime:** A spinner box set to '1' with '(1 to 24 hours)' in parentheses.
- Maximum Cache Entries:** A spinner box set to '128' with '(10 to 1,024)' in parentheses.

 At the bottom right are 'OK', 'Reset', and 'Exit' buttons.

FIGURE 343 RADIUS Server Policy screen - Server Policy tab

8. The **Server Policy** tab displays by default.
9. If creating a new policy, assign it a **RADIUS Server Policy** name up to 32 characters.
10. Set the following **Settings** required in the creation or modification of the server policy:.

RADIUS User Pools

Select the user pools to apply to this server policy. Up to 32 policies can be applied.

LDAP Server Dead Period

Set an interval in either *Seconds* (0 - 600) or *Minutes* (0- 10) during which the controller will not contact its LDAP server resource. A dead period is only implemented when additional LDAP servers are configured and available.

LDAP Groups

Use the drop-down menu to select LDAP groups to apply the server policy configuration. Select the Create or Edit icons as needed to either create a new group or modify an existing group. Use the arrow icons to add and remove groups as required.

LDAP Group Verification

Select the checkbox to set the LDAP group search configuration.

Local Realm

Define the LDAP performing authentication using information from an LDAP server. User information includes user name, password, and the groups to which the user belongs.

11. Set the following **Authentication** parameters to define server policy authorization settings.

Authentication Data Source	Select the RADIUS resource for user authentication with this server policy. Options include <i>Local</i> for the controller's local user database or <i>LDAP</i> for a remote LDAP resource. The default setting is Local.
Local Authentication Type	Use the drop-down menu to select the controllers local EAP authentication scheme used with this policy. The following EAP authentication types are supported by the controller's onboard RADIUS server: <i>All</i> - Enables both TTLS and PEAP. <i>TLS</i> - Uses TLS as the EAP type <i>TLS and MD5</i> - The EAP type is TTLS with default authentication using MD5. <i>TTLS and PAP</i> - The EAP type is TTLS with default authentication using PAP. <i>TTLS and MSCHAPv2</i> - The EAP type is TTLS with default authentication using MSCHAPv2. <i>PEAP and GTC</i> - The EAP type is PEAP with default authentication using GTC. <i>PEAP and MSCHAPv2</i> - The EAP type is PEAP with default authentication using MSCHAPv2.
LDAP Authentication Type	Use the drop-down menu to select the LDAP authentication scheme used with this policy. The following LDAP authentication types are supported by the controller's external LDAP resource: <i>All</i> - Enables both TTLS and PAP and PEAP and GTC. <i>TTLS and PAP</i> - The EAP type is TTLS with default authentication using PAP. <i>PEAP and GTC</i> - The EAP type is PEAP with default authentication using GTC.
Enable CRL Validation	Select this option to enable a <i>Certificate Revocation List</i> (CRL) check is made. Certificates can be checked and revoked for a number of reasons including failure or compromise of a device using a certificate, a compromise of a certificate key pair or errors within an issued certificate. This option is disabled by default.

12. Set the following **Session Resumption/Fast Reauthentication** settings to define how server policy sessions are re-established once terminated and require cached data to resume:

Enable Session Resumption	Select the checkbox to control volume and the duration cached data is maintained by the server policy upon the termination of a server policy session. The availability and quick retrieval of the cached data speeds up session resumption.
Cached Entry Lifetime	Use the spinner control to set the lifetime (1 - 24 hours) cached data is maintained by the RADIUS server policy. The default setting is 1 hour.
Maximum Cache Entries	Use the spinner control to define the maximum number of entries maintained in cache for this RADIUS server policy. The default setting is 128 entries.

13. Select **OK** to save the settings to the server policy configuration. Select **Reset** to revert to the last saved configuration.

14. Refer to the following to add RADIUS clients, proxy server configurations, LDAP server configurations and review deployment considerations impacting the effectiveness of the controller's RADIUS deployment:

- [Configuring RADIUS Clients](#)
- [Configuring a RADIUS Proxy](#)
- [Configuring an LDAP Server Configuration](#)

Configuring RADIUS Clients

The controller uses a RADIUS client as a mechanism to communicate with a central server to authenticate users and authorize access to the controller managed network.

The client and server share a secret. That shared secret followed by the request authenticator is put through a MD5 hash to create a 16 octet value which is XORed with the password entered by the user. If the user password is greater than 16 octets, additional MD5 calculations are performed, using the previous ciphertext instead of the request authenticator. The server receives a RADIUS *access request* packet and verifies the server possesses a shared secret for the client. If the server does not possess a shared secret for the client, the request is dropped. If the client received a verified *access accept* packet, the username and password are considered correct, and the user is authenticated. If the client receives a verified *access reject* message, the username and password are considered to be incorrect, and the user is not authenticated.

1. Select the **Client** tab from the RADIUS Server Policy screen.

The screenshot shows the 'RADIUS Server Policy' interface for 'policy1'. It has tabs for 'Server Policy', 'Client', 'Proxy', and 'LDAP'. The 'Client' tab is active, showing a table of RADIUS Clients. The table has two columns: 'IP Address' and 'Shared Secret'. The first row shows an IP address of '15.125.81.5/24' and a masked shared secret. There are edit, delete, and show/hide icons for each row. An 'Add Row' button is located below the table. At the bottom of the window are 'OK', 'Reset', and 'Exit' buttons.

IP Address	Shared Secret	
15.125.81.5/24	*****	Show

FIGURE 344 RADIUS Server Policy screen - Client tab

2. Select the **+ Add Row** button to add a table entry for a new client's IP address, mask and shared secret. To delete a client entry, select the **Delete** icon on the right-hand side of the table entry.
3. Specify the **IP Address** and mask of the RADIUS client authenticating with the controller managed RADIUS server.
4. Specify a **Shared Secret** for authenticating the RADIUS client.

5. Shared secrets verify RADIUS messages with RADIUS enabled device configured with the same shared secret. Select the **Show** checkbox to expose the shared secret's actual character string, leaving the option unselected displays the shared secret as a string of asterisks (*).
6. Click **OK** button to save the server policy's client configuration. Click the **Reset** button to revert to the last saved configuration.

Configuring a RADIUS Proxy

A user's access request is sent to a proxy server if it cannot be authenticated by the controller's local RADIUS resources. The proxy server checks the information in the user access request and either accepts or rejects the request. If the proxy server accepts the request, it returns configuration information specifying the type of connection service required to authenticate the user.

The RADIUS proxy appears to act as a RADIUS server to the NAS, whereas the proxy appears to act as a RADIUS client to the RADIUS server.

When the controller's RADIUS server receives a request for a user name containing a realm, the server references a table of configured realms. If the realm is known, the server proxies the request to the RADIUS server. The behavior of the proxying server is configuration-dependent on most servers. In addition, the proxying server can be configured to add, remove or rewrite requests when they are proxied.

To define a proxy configuration:

1. Select the **Proxy** tab from the RADIUS Server Policy screen.

The screenshot shows the 'RADIUS Server Policy' configuration window for 'policy1'. It has tabs for 'Server Policy', 'Client', 'Proxy', and 'LDAP'. The 'Proxy' tab is active.

Proxy Retries

Proxy Retry Delay: 5 seconds (5 to 10)

Proxy Retry Count: 3 (3 to 6)

Realms

Realm Name	IP Address	Port Number	Shared Secret
engineering 1	212 . 4 . 56 . 9	1812	*****

Buttons: Add Row, OK, Reset, Exit

FIGURE 345 RADIUS Server Policy screen - Proxy tab

2. Enter the Proxy server retry delay time in the **Proxy Retry Delay** field. Enter a value in seconds within the range of 5-10 seconds. This is the interval the controller's RADIUS server waits before making an additional connection attempt. The default delay interval is 5 seconds.
3. Enter the Proxy server retry count value in the **Proxy Retry Count** field. Enter a value within the range of 3-6 to define the number of retries sent to proxy server before giving up the request. The default retry count is 3 attempts.
4. Select the **+ Add Row** button to add a RADIUS server proxy realm name and network address. To delete a proxy server entry, select the **Delete** icon on the right-hand side of the table entry.
5. Enter the realm name in the **Realm Name** field. The realm name cannot exceed 50 characters. When the controller's RADIUS server receives a request for a user name with a realm, the server references a table of realms. If the realm is known, the server proxies the request to the RADIUS server.
6. Enter the Proxy server IP address in the **IP Address** field. This is the address of server checking the information in the user access request and either accepting or rejecting the request on behalf of the controller's RADIUS server.
7. Enter the TCP/IP port number for the server that acts as a data source for the proxy server in the **Port Number** field. Use the spinner to select a value between 1024-65535. The default port is 1812.
8. Enter the RADIUS client shared secret in the **Shared Secret** field for authenticating the RADIUS proxy.

2. Refer to the following to determine whether an LDAP server can be used as is, a server configuration requires creation or modification or a configuration requires deletion and permanent removal.

Redundancy

Displays whether the listed LDAP server IP address has been defined as a primary or secondary server resource. Designating at least one secondary server is a good practice to ensure RADIUS user information is available if a primary server were to become unavailable.

IP Address

Displays the IP address of the external LDAP server acting as the data source for the controller's RADIUS server.

Port

Lists the physical port number used by the controller's RADIUS server to secure a connection with the remote LDAP server resource.

3. Click the **Add** button to add a new LDAP server configuration, **Edit** to modify an existing LDAP server configuration or **Delete** to remove a LDAP server from the list of those available.

FIGURE 347 LDAP Server Add screen

4. Set the following **Network** address information required for the connection to the external LDAP server resource:.

Redundancy

Define whether this LDAP server is a primary or secondary server resource. Primary servers are always queried for connection first. However, designating at least one secondary server is a good practice to ensure RADIUS user information is available if a primary server were to become unavailable.

IP Address

Set the IP address of the external LDAP server acting as the data source for the controller's RADIUS server.

Login

Define a unique login name used for accessing the controller's remote LDAP server resource. Consider using a unique login name for each LDAP server provided to increase the security of the connection between the controller and remote LDAP resource.

Port

Use the spinner control to set the physical port number used by the controller's RADIUS server to secure a connection with the remote LDAP server resource.

Timeout

Set an interval between 1 - 10 seconds the controller's RADIUS server uses as a wait period for a response from the target primary or secondary LDAP server resource. The default setting is 10 seconds.

- Set the following **Network** address information required for the connection to the external LDAP server resource:

Bind DN	Specify the distinguished name to bind with the LDAP server. The DN is the name that uniquely identifies an entry in the LDAP directory. A DN is made up of attribute value pairs, separated by commas.
Base DN	Specify a <i>distinguished name</i> (DN) that establishes the base object for the search. The base object is the point in the LDAP tree at which to start searching. LDAP DNs begin with the most specific attribute (usually some sort of name), and continue with progressively broader attributes, often ending with a country attribute. The first component of the DN is referred to as the <i>Relative Distinguished Name</i> (RDN). It identifies an entry distinctly from any other entries that have the same parent.
Bind Password	Enter a valid password for the LDAP server. Select the Show checkbox to expose the password's actual character string, leaving the option unselected displays the password as a string of asterisks (*). The password cannot 32 characters.
Password Attribute	Enter the LDAP server password attribute. The password cannot exceed 64 characters.

- Set the following **Attributes** for LDAP groups to optimally refine group queries:

Group Attribute	LDAP systems have the facility to poll dynamic groups. In an LDAP dynamic group an administrator can specify search criteria. All users matching the search criteria are considered a member of this dynamic group. Specify a group attribute used by the LDAP server. An attribute could be a group name, group ID, password or group membership name.
Group Filter	Specify the group filters used by the LDAP server. This filter is typically used for security role-to-group assignments and specifies the property to look up groups in the directory service.
Group Membership Attribute	Specify the group member attribute sent to the LDAP server when authenticating users.

- Click the **OK** button to save the changes to the LDAP server configuration. Select **Reset** to revert to the last saved configuration.

RADIUS Deployment Considerations

Before defining the controller's internal RADIUS server configuration, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- Brocade Mobility recommends each RADIUS client use a different shared secret. If a shared secret is compromised, only the one client poses a risk to the data managed by the controller as opposed all the additional clients that potentially share the secret password.
- Consider using an LDAP server as a database of user credentials that can be used optionally with the controller's RADIUS server to free up resources and manage user credentials from a secure remote location.

Management Access

In this chapter

- [Viewing Management Access Policies](#) 511
- [Management Access Deployment Considerations](#) 525

The controller has mechanisms to allow/deny Management Access for separate interfaces and protocols (HTTP, HTTPS, Telnet, SSH or SNMP). Management access can be enabled/disabled as required for unique policies. The controller's Management Access functionality is not meant to function as an ACL (in routers or other firewalls), where administrators specify and customize specific IPs to access specific interfaces.

Brocade Mobility recommends disabling un-used and insecure interfaces as required within managed access profiles. Disabling un-used management services can dramatically reduce an attack footprint and free resources on devices managed by the controller.

Viewing Management Access Policies

Brocade Mobility controllers can be remotely managed using several interfaces including SNMP, CLI and HTTP/HTTPS.

Controller Management Access policies display in the lower left-hand side of the screen. Existing policies can be updated as controller management permissions change, or new policies can be added as needed.

To view existing Management Access policies:

1. Select **Configuration > Management > Management Policy** to display the main Management Policy screen and Management Browser.
2. Select a policy from the Management Browser or refer to the Management screen (displayed by default) to review existing Management Access policy configurations at a higher level.

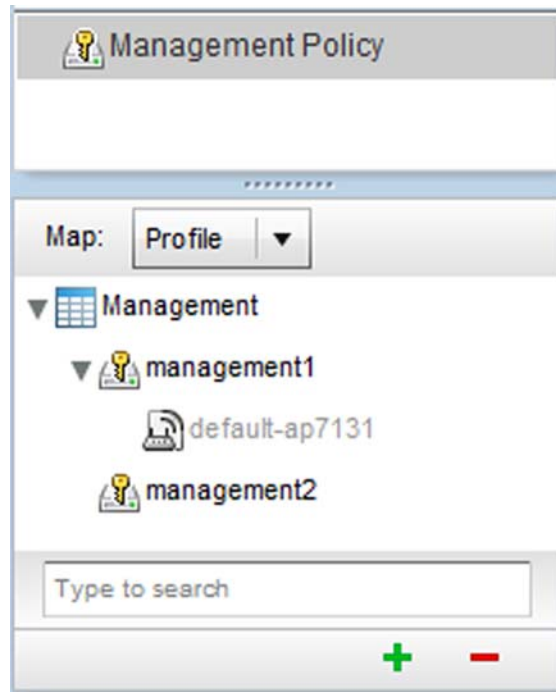


FIGURE 348 Management Browser screen

3. The **Management Policy** screen displays existing management policies and their unique protocol support configurations.

The screenshot shows the 'Management Policy' screen with the following table content:

Management Policy	Telnet	SSHv2	HTTP	HTTPS	SHMPv2	SHMPv3	FTP
management1	✓	✓	✗	✓	✗	✗	✓
management2	✗	✓	✓	✗	✓	✓	✗

At the bottom of the table, there is a search bar labeled 'Type to search in tables', a 'Row Count: 2' indicator, and three buttons: 'Add', 'Edit', and 'Delete'.

FIGURE 349 Management Policy screen

- Refer to the following Management Access policy configurations to discern whether these existing policies can be used as is, require modification or a new policy requires creation:

A green check mark indicates controller device access is allowed using the listed protocol. A red X indicates controller device access is denied using the listed protocol.

- Management Policy** Displays the name of the Management Access policy assigned when initially created. The name cannot be updated when modifying a policy.
- Telnet** Telnet provides a command line interface to a remote host over TCP. Telnet provides no encryption, but it does provide a measure of authentication.
- SSHv2** SSH (*Secure Shell*) version 2, like Telnet, provides a command line interface to a remote host. However, all SSH transmissions are encrypted, increasing their security.
- HTTP** HTTP (*Hypertext Transfer Protocol*) provides access to the device’s GUI using a Web browser. This protocol is not very secure.
- HTTPS** HTTPS (*Hypertext Transfer Protocol Secure*) provides fairly secure access to the device’s GUI using a Web browser. Unlike HTTP, HTTPS uses encryption for transmission, and is therefore more secure.

SNMPv2	SNMP (<i>Simple Network Management Protocol</i>) exposes a device's management data so it can be managed remotely. Device data is exposed as variables that can be accessed and modified. However, SNMP is generally used to monitor a system's performance and other parameters.
SNMPv3	SNMP (<i>Simple Network Management Protocol</i>) exposes a device's management data so it can be managed remotely. Device data is exposed as variables that can be accessed and modified. However, SNMP is generally used to monitor a system's performance and other parameters.
FTP	FTP (<i>File Transfer Protocol</i>) is a standard protocol for files transfers over a TCP/IP network.

5. If it's determined a Management Access policy requires creation or modification, refer to *Adding or Editing a Management Access Policy on page 11-514*. If necessary, select an existing Management Access policy and select Delete to permanently remove it from the list of those available.

Adding or Editing a Management Access Policy

[Viewing Management Access Policies](#)

To add a new Management Access policy, or edit an existing configuration:

1. Select **Configuration > Management > Wireless LAN Policy** to the main Management Policy screen and Management Browser.
2. Existing policies can be modified by either selecting a policy from the **Management Browser** and selecting the **Edit** button.
3. New policies can be created by selecting the **Add** button from the bottom right-hand side of the Management screen.
4. A name must be supplied to the new policy before the **Access Control**, **SNMP**, **SNMP Traps** and **Administrators** tabs become enabled and the policy's configuration defined. The name cannot exceed 32 characters.




FIGURE 350 Management Policy screen - New Policy Creation

5. Select **OK** to commit the new policy name.

Once the new name is defined, the screen's four tabs become enabled, with the contents of the **Administrators** tab displayed by default. Refer to the following to define the configuration of the new controller Management Access policy:

- [Creating an Administrator Configuration](#) - Use the *Administrators* tab to create specific controller users, assign them permissions to specific protocols and set specific administrative roles for the managed network.
- [Setting the Access Control Configuration](#) - Use the *Access Control* tab to enable/disable specific protocols and interfaces. Again, this kind of access control is not meant to function as an ACL, but rather as a means to enable/disable specific protocols (HTTP, HTTPS, Telnet etc.) for each Management Access policy.
- [Setting the Authentication Configuration](#) - Refer to the *Authentication* tab to set the authentication scheme used to validate user credentials with this policy.
- [Setting the SNMP Configuration](#) - Refer to the *SNMP* tab to enable SNMPv2, SNMPv3 or both and define specific community strings for this policy.
- [SNMP Trap Configuration](#) - Use the *SNMP Traps* tab to enable trap generation for the policy and define trap receiver configurations.

For deployment considerations and recommendations impacting a controller's Management Access policy configuration, refer to *Management Access Deployment Considerations on page 11-525*.

The screenshot shows the 'Administrators' configuration screen. At the top, the 'User Name' field contains 'installer'. Below this, the 'Settings' section includes a 'Password' field (masked with asterisks) and a 'Reconfirm' field (also masked). The 'Access' section features an 'Access Type' label and four checked checkboxes: 'Web UI', 'Telnet', 'SSH', and 'Console'. The 'Administrator Roles' section has a 'Role' label and seven radio button options: 'Superuser', 'System', 'Network' (which is selected), 'Security', 'Monitor', 'Help Desk', and 'Web User'. At the bottom right, there are three buttons: '> OK', 'Reset', and 'Exit'.

FIGURE 352 Administrators screen

4. If creating a new administrator, enter a user name in the **User Name** field. This is a mandatory field for new administrators and cannot exceed 32 characters. Optimally assign a name representative of the user and role.
5. Provide a strong password for the administrator in the **Password** field, once provided, **Reconfirm** the password to ensure its accurately entered. This is a mandatory field.
6. Select **Access** options to define the permitted access for the user. If required, all four options can be selected and invoked simultaneously.

Web UI	Select this option to enable access to the device's Web User Interface.
Telnet	Select this option to enable access to the device using TELNET.
SSH	Select this option to enable access to the device using SSH.
Console	Select this option to enable access to the device's console.

7. Select the **Administrator Role** for the administrator using this profile. Only one role can be assigned.

Superuser	Select this option to assign complete administrative rights to the user. This entails all the roles listed for all the other administrative roles.
System	Select System to allow the administrator to configure general settings like NTP, boot parameters, licenses, perform image upgrade, auto install, manager redundancy/clustering and control access.
Network	Select this option to allow the user to configure all wired and wireless parameters (IP configuration, VLANs, L2/L3 security, WLANs, radios etc).
Security	Select Security to set the administrative rights for a security administrator allowing configuration of all security parameters.

Monitor	Select Monitor to assign permissions without any administrative rights. The Monitor option provides read-only permissions.
Help Desk	Assign this role to someone who typically troubleshoots and debugs problems reported by the customer. The Help Desk manager typically runs troubleshooting utilities (like a sniffer), executes service commands, views/retrieves logs and reboots the controller.
Web User	Select Web User to assign the administrator privileges needed to add users for authentication.

8. Select the **OK** button to save the administrator's configuration. Select **Reset** to revert to the last saved configuration.

Setting the Access Control Configuration

Adding or Editing a Management Access Policy

Refer to the **Access Control** tab to allow/deny management access to the managed network using strategically selected protocols (HTTP, HTTPS, Telnet, SSH or SNMP). Access options can be either enabled or disabled as required. Brocade Mobility recommends disabling unused interfaces to close unnecessary security holes. The Access Control tab is not meant to function as an ACL (in routers or other firewalls), where you can specify and customize specific IPs to access specific interfaces.

The following table demonstrates some interfaces provide better security than others:

Access Type	Encrypted	Authenticated	Default State
Telnet	No	Yes	Disabled
HTTP	No	Yes	Disabled
HTTPS	Yes	Yes	Disabled
SSHv2	Yes	Yes	Disabled

To set an access control configuration for the Management Access policy:

1. Select the **Access Control** tab from the Management Policy screen.

The screenshot shows the 'Management Policy' configuration window for 'admin1'. The 'Access Control' tab is selected. The configuration is organized into several sections:

- Telnet:** 'Enable Telnet' is unchecked. 'Telnet Port' is set to 23 (range 1 to 65,535).
- SSH:** 'Enable SSHv2' is checked. 'SSHv2 Port' is set to 22.
- HTTP/HTTPS:** 'Enable HTTP' is checked. 'Enable HTTPS' is unchecked.
- FTP:** 'Enable FTP' is unchecked. 'FTP Username' is 'ftuser', 'FTP Password' is empty, and 'FTP Root Directory' is 'flash:/'.
- General:** 'Idle Session Timeout' is 120 (range 0 to 1,440). 'Message of the Day' is empty.
- Access Restrictions:** 'Filter Type' is empty. 'IP Access List' is '<none>'. 'Source Hosts' shows a list of IP addresses (0.0.0.0) with 'Clear' buttons. 'Source Subnets' is empty. 'Logging Policy' is empty.

Buttons for 'OK', 'Reset', and 'Exit' are at the bottom right.

FIGURE 353 Management Policy screen - Access Control tab

2. Set the following parameters required for **Telnet** access:

Enable Telnet

Select the checkbox to enable Telnet device access. Telnet provides a command line interface to a remote host over TCP. Telnet provides no encryption, but it does provide a measure of authentication. Telnet access is disabled by default.

Telnet Port

Set the port on which Telnet connections are made (1 - 65,535). The default port is 23. Change this value using the spinner control next to this field or by entering the port number in the field.

3. Set the following parameters required for **SSH** access:

Enable SSHv2

Select the checkbox to enable SSH device access. SSH (*Secure Shell*) version 2, like Telnet, provides a command line interface to a remote host. SSH transmissions are encrypted and authenticated, increasing the security of transmission. SSH access is disabled by default.

SSHv2 Port

Set the port on which SSH connections are made. The default port is 22. Change this value using the spinner control next to this field or by entering the port number in the field.

4. Set the following **HTTP/HTTPS** parameters:

Enable HTTP	Select the checkbox to enable HTTP device access. HTTP provides limited authentication and no encryption.
Enable HTTPS	Select the checkbox to enable HTTPS device access. HTTPS (<i>Hypertext Transfer Protocol Secure</i>) is more secure plain HTTP. HTTPS provides both authentication and data encryption as opposed to just authentication.

NOTE

If the a RADIUS server is not reachable, HTTPS or SSH management access to the controller or Access Point may be denied.

5. Set the following parameters required for **FTP** access:

Enable FTP	Select the checkbox to enable FTP device access. FTP (<i>File Transfer Protocol</i>) is the standard protocol for transferring files over a TCP/IP network. FTP requires administrators enter a valid username and password authenticated locally on the controller. FTP access is disabled by default.
FTP Username	Specify a username required when logging in to the FTP server. The username cannot exceed 32 characters.
FTP Password	Specify a password required when logging in to the FTP server. Reconfirm the password in the field provided to ensure it has been entered correctly. The password cannot exceed 63 characters.
FTP Root Directory	Provide the complete path to the root directory in the space provided. The default setting has the root directory set to flash:/

6. Set the following **General** parameters::

Idle Session Timeout	Specify a inactivity timeout for management connects (in seconds) between 0 - 1,440.
Message of the Day	Enter message of the day text to be displayed at login for clients connecting via Telnet or SSH.

7. Set the following **Access Restriction** parameters::

Filter Type	Select a filter type for access restriction, either ip-access-list, source-address, or none.
IP Access List	If the selected filter type is ip-access-list, select an ip access list from the drop-down menu or select the create button to make a new one.
Source Hosts	If the selected filter type is source-address, enter an IP Address or IP Addresses for the source hosts.
Source Subnets	If the selected filter type is source-address, enter a source subnet or subnets for the source hosts.
Logging Policy	If the selected filter type is source-address, enter a logging policy as none, denied-only or All.

8. Select **OK** to update the access control configuration. Select **Reset** to the last saved configuration.

Setting the Authentication Configuration

Adding or Editing a Management Access Policy

Refer to the **Authentication** tab to define how user credential validation is conducted on behalf of a Management Access policy

To configure an external authentication resource:

1. Select the **Authentication** tab from the Management Policy screen.

The screenshot shows the 'Management Policy' configuration window for 'admin1'. The 'Authentication' tab is selected. The configuration options are as follows:

- Local:** Radio buttons for 'Enabled' (selected) and 'Disabled'.
- RADIUS:** Checkboxes for 'External' and 'Fallback'.
- IP Address:** A text input field containing '0 . 0 . 0 . 0'.
- UDP Port:** A spinner control set to '1812'.
- Shared Secret:** An empty text input field.
- Attempts:** A spinner control set to '3' with a range of '(1 to 10)'.
- Timeout:** A spinner control set to '3' with a unit dropdown set to 'Seconds' and a range of '(1 to 60)'.

At the bottom right, there are buttons for 'OK', 'Reset', and 'Exit'.

FIGURE 354 Management Policy screen - Authentication tab

2. Define the following settings to authenticate management access requests:

Local	Select whether the Access Point's authentication server resource is centralized (local) to the Access Point itself, or whether the Access Point shall use an external authentication resource for validating user access.
RADIUS	If local authentication is disabled, set whether the RADIUS server type is External and or Fallback.
IP Address	Define the numerical IP address of the Access Point's external RADIUS authentication resource.
UDP Port	Use the spinner control to set the port number where the RADIUS server is listening. The default setting is 1812.

Shared Secret	Define a shared secret password between the Access Point and the RADIUS server that must be provided to secure the external RADIUS resource.
Attempts	Set the number of times an authentication request is sent to the RADIUS server before giving up. The available range is 1- 10, with a default of 3.
Timeout	Set a timeout setting in Seconds (1-60) after which requests to the RADIUS server will be retries.

3. Select **OK** to update the authentication configuration. Select **Reset** to the last saved configuration.

Setting the SNMP Configuration

Adding or Editing a Management Access Policy

The controller can use the *Simple Network Management Protocol* (SNMP) to communicate with devices within the managed network. SNMP is an application layer protocol that facilitates the exchange of management information between the controller and a managed device. SNMP enabled devices listen on port 162 (by default) for SNMP packets from the controller's management server. SNMP uses read-only and read-write community strings as an authentication mechanism to monitor and configure supported devices. The read-only community string is used to gather statistics and configuration parameters from a supported wireless device. The read-write community string is used by a management server to set device parameters. SNMP is generally used to monitor a system's performance and other parameters.

SNMP Version	Encrypted	Authenticated	Default State
SNMPv2	No	No	Enabled
SNMPv3	Yes	Yes	Enabled

To configure SNMP Management Access within the managed network:

1. Select the **SNMP** tab from the Management Policy screen.

Management Policy admin1

Administrators Access Control Authentication SHMP SHMP Traps

SNMP

Enable SNMPv2

Enable SNMPv3

SNMP v1/v2c Community String

Community	Access Control	

Add Row

SNMPv3 Users

User Name	Authentication	Encryption	Password	

Add Row

OK Reset Exit

FIGURE 355 Management Policy screen - SNMP tab

2. Enable or disable SNMPv2 and SNMPv3.

Enable SNMPv2

Select the checkbox to enable SNMPv2 support. SNMPv2 provides device management using a hierarchical set of variables. SNMPv2 uses *Get*, *GetNext*, and *Set* operations for data management. SNMPv2 is enabled by default.

Enable SNMPv3

Select the checkbox to enable SNMPv3 support. SNMPv3 adds security and remote configuration capabilities to previous versions. The SNMPv3 architecture introduces the *User-based Security Model (USM)* for message security and the *View-based Access Control Model (VACM)* for access control. The architecture supports the concurrent use of different security, access control and message processing techniques. SNMPv3 is enabled by default.

3. Set the **SNMP v1/v2 Community String** configuration. Use the **+ Add Row** function as needed to add additional SNMP v1/2 community strings, or select an existing community string's radio button and select the **Delete** icon to remove it.

Community

Define a *public* or *private* community designation. By default, SNMPv2 community strings on most devices are set to *public* for the read-only community string and *private* for the read-write community string.

Access Control

Set the access permission for each community string used by devices to retrieve or modify information. The available options are:

Read Only - Allows a remote device to retrieve information

Read-Write - Allows a remote device to modify settings

4. Set the **SNMPv3 Users** configuration. Use the **+ Add Row** function as needed to add additional SNMP v3 user configurations, or select a SNMP user's radio button and select the **Delete** icon to remove the user.

User Name	Use the drop-down menu to define a user name of <i>snmpmanager</i> , <i>snmpoperator</i> or <i>snmptrap</i> .
Authentication	Displays the authentication scheme used with the listed SNMPv3 user. The listed authentication scheme ensures only trusted and authorized users and devices can access the managed network.
Encryption	Displays the encryption scheme used with the listed SNMPv3 user.
Password	Provide the user's password in the field provided. Select the Show check box to display the actual character string used in the password, while leaving the check box unselected protects the password and displays each character as "*".

5. Select **OK** to update the SNMP configuration. Select **Reset** to revert to the last saved configuration.

SNMP Trap Configuration

Adding or Editing a Management Access Policy

The managed network can use SNMP trap receivers for fault notifications. SNMP traps are unsolicited notifications triggered by thresholds (or actions), and are therefore an important fault management tool.

A SNMP trap receiver is the destination of SNMP messages (external to the controller). A trap is like a Syslog message, just over another protocol (SNMP). A trap is generated when a device consolidates event information and transmits the information to an external repository. The trap contains several standard items, such as the SNMP version, community etc.

SNMP trap notifications exist for most controller operations, but not all are necessary for day-to-day operation.

To define a SNMP trap configuration for receiving events at a remote destination:

1. Select the **SNMP Traps** tab from the Management Policy screen.

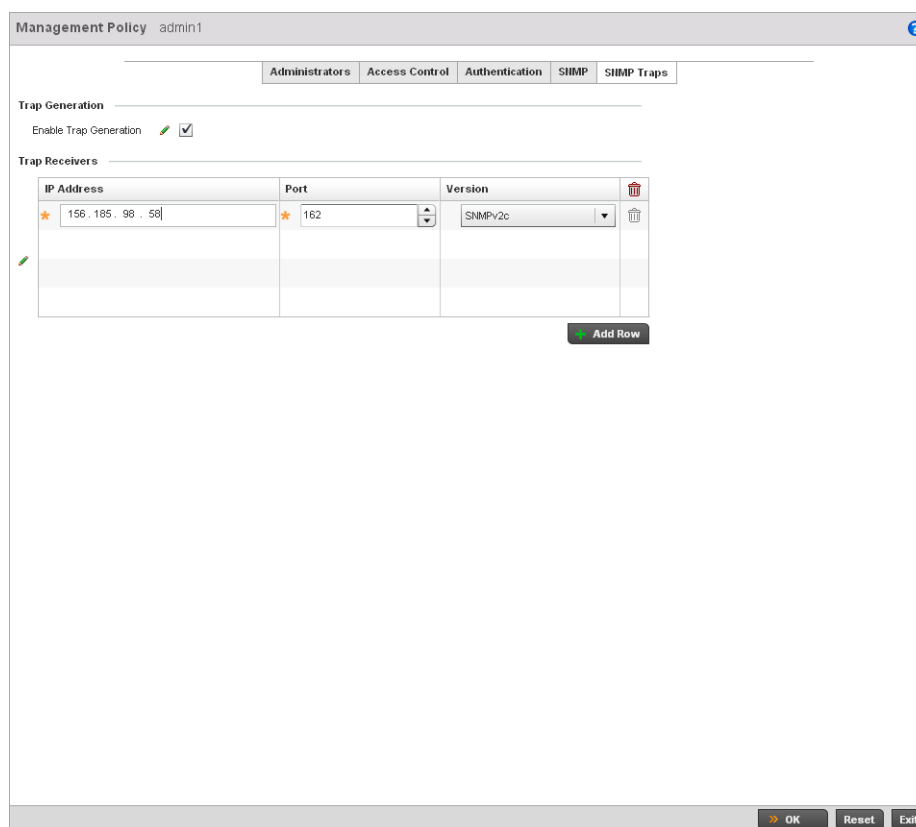


FIGURE 356 Management Policy screen - SNMP Traps tab

2. Select the **Enable Trap Generation** checkbox to enable trap generation using the trap receiver configuration defined. This feature is disabled by default.
3. Refer to the **Trap Receiver** table to set the configuration of the external resource dedicated to receiving trap information on behalf of the controller. Select **Add Row +** as needed to add additional trap receivers. Select the **Delete** icon to permanently remove a trap receiver.

IP Address Sets the IP address of the external server resource dedicated to receiving the SNMP traps on behalf of the controller.

Port Set the port of the server resource dedicated to receiving SNMP traps. The default port is port 162.

Version Sets the SNMP version to use to send SNMP traps. SNMPv2 is the default.

4. Select **OK** to update the SNMP Trap configuration. Select **Reset** to revert to the last saved configuration.

Management Access Deployment Considerations

Before defining a access control configuration as part of a Management Access policy, refer to the following deployment guidelines to ensure the configuration is optimally effective:

- Unused management protocols should be disabled to reduce a potential attack against managed resources.
- Use management interfaces providing encryption and authentication. Management services like HTTPS, SSH and SNMPv3 should be used when possible, as they provide both data privacy and authentication.
- By default, SNMPv2 community strings on most devices are set to *public* for the read-only community string and *private* for the read-write community string. Legacy Brocade Mobility devices may use other community strings by default.
- Brocade Mobility recommends SNMPv3 be used for device management, as it provides both encryption, and authentication.
- Enabling SNMP traps can provide alerts for isolated attacks at both small managed radio deployments or distributed attacks occurring across multiple managed sites.
- Whenever possible, Brocade Mobility recommends centralized RADIUS management be enabled on controllers and Access Points. This provides better management and control of management usernames and passwords and allows administrators to quickly change credentials in the event of a security breach.

Diagnostics

In this chapter

- [Fault Management](#) 527
- [Snapshots](#) 532
- [Crash Files](#) 534
- [Advanced Diagnostics](#)..... 536

The controller's resident diagnostic capabilities enable administrators to understand how managed devices are performing and troubleshoot issues impacting network performance. Performance and diagnostic information is collected and measured on Brocade Mobility controllers and Access Points for any anomalies causing a key controller processes to potentially fail.

Numerous tools are available within the Diagnostics menu. Some allow event filtering, some allow you to view logs and some allow you to manage files generated when major hardware or software issues are detected.

Fault Management

Fault management enables user's administering multiple sites to assess how individual devices are performing and review issues impacting the managed network. Use the controller's Fault Management screens to administrate errors generated by the controller, Access Point or wireless clients managed by the controller.

1. Select **Diagnostics > Fault Management**.
2. The **Filter Events** screen displays by default. Use this screen to configure how events are tracked. By default, all events are enabled, and an administrator has to turn off events that do not require tracking.

FIGURE 357 Fault Management Filter Events screen

3. Use the **Filter Events** screen to create filters for managing displayed events. Events can be filtered based on severity, the module received, source MAC, device MAC and client MAC address.
4. Define the following **Customize Event Filters** parameters for the Fault Management configuration:

Severity

Set the filtering severity. Select from the following:

All Severities – All events are displayed irrespective of their severity

Critical – Only critical events are displayed

Error – Only errors and above are displayed

Warning – Only warnings and above are displayed

Informational – Only informational and above events are displayed

Module

Select the module from which events are tracked. When a module is selected, events from other modules are not tracked. Remember this when interested in events generated by a particular controller module. Individual modules can be selected (such as TEST, LOG, FSM etc.) or all modules can be tracked by selecting All Modules.

Source

Set the MAC address of the source device to be tracked. Setting a MAC address of 00:00:00:00:00:00 allows all devices to be tracked.

Device

Set the device MAC address for the device (such as an Access Point or wireless client) from which the source MAC address is tracked. Setting a MAC address of 00:00:00:00:00:00 allows all devices.

8. Use the **View Events** screen to track and troubleshoot events using source and severity levels defined in the Configure events screen.
9. Define the following **Customize Event Filters** parameters for the Fault Management configuration:

Timestamp	Displays the Timestamp (time zone specific) when the fault occurred.
Module	Displays the module used to track the event. Events detected by other module are not tracked.
Message	Displays error or status messages for each event listed.
Severity	Displays the severity of the event as defined for tracking from the Configuration screen. Severity options include: <i>All Severities</i> – All events are displayed irrespective of their severity <i>Critical</i> – Only critical events are displayed <i>Error</i> – Only errors and above are displayed <i>Warning</i> – Only warnings and above are displayed <i>Informational</i> – Only informational and above events are displayed
Source	Displays the MAC address of the source device tracked by the controller and selected module.

10. Select **Clear All** to clear the events displayed and begin new event data gathering.
11. Select **Event History** from the upper, left-hand, side of the **Diagnostics > Fault Management** menu.

Severity	Displays the severity of the event as defined for tracking from the Configuration screen. Severity options include: <i>All Severities</i> – All events are displayed irrespective of their severity <i>Critical</i> – Only critical events are displayed <i>Error</i> – Only errors and above are displayed <i>Warning</i> – Only warnings and above are displayed <i>Informational</i> – Only informational and above events are displayed
Source	Displays the MAC address of the source device tracked by the controller and selected module.
RF Domain	Displays the RF Domain of the source device tracked by the controller and selected module.

15. Clicking the **Fetch Historical Events** button retrieves all log history for the device in the Select Device drop-down.
16. Select **Clear All** to clear the events displayed and begin new event data gathering.

Snapshots

Use the Snapshots screens to review panic and core dump files created when a controller or managed device encounters a critical error or malfunction.

Core Snapshots

[Snapshots](#)

Refer to the **Core Snapshots** screen to review core dump files (system events and process failures with a.core extension) and troubleshoot issues specific to the device on which the event was generated. Core snapshots are issues impacting the controller core (distribution layer). Once reviewed, core files can be deleted or transferred for archive. Core files can be sent to the Brocade Mobility support team to expedite any issues that arise with the reporting device.

To review core snapshots impacting the managed network:

1. Select **Diagnostics > Snapshots**.
 The Core Snapshots screen displays by default. This screen displays a list of device MAC addresses impacted by core dumps.
2. Select a device from those displayed in the lower, left-hand, side of the controller UI.

1. Select **Diagnostics > Snapshots**.
2. Select **Panic Snapshots** from the upper, left-hand, side of the controller UI. A list of device MAC addresses impacted by panic events displays.
3. Select a device from those displayed in the lower, left-hand, side of the controller UI.

Device	System Name	Type
AA-00-00-00-00-00	ControllerA-RFS7000	RFS7000
AA-11-00-00-00-00	AP1-ControllerA-AP650	AP650
AA-22-00-00-00-00	AP2-ControllerA-AP71xx	AP71xx
AA-33-00-00-00-00	AP3-ControllerA-AP71xx	AP71xx
AA-44-00-00-00-00	AP4-ControllerA-AP71xx	AP71xx
BB-00-00-00-00-00-00	ControllerB-RFS6000	RFS6000
BB-11-00-00-00-00	AP1-ControllerB-AP650	AP650
BB-22-00-00-00-00	AP2-ControllerB-AP650	AP650
BB-33-00-00-00-00	AP3-ControllerB-AP650	AP650
BB-44-00-00-00-00	AP4-ControllerB-AP6511	AP6511
BB-55-00-00-00-00	AP5-ControllerB-AP621	ap621
BB-66-00-00-00-00	AP6-ControllerB-AP621	ap621
BB-77-00-00-00-00	AP7-ControllerB-AP632	AP6532
CC-00-00-00-00-00	ControllerC-RFS4000	RFS4000
DD-00-00-00-00-00	ControllerD-RFS4000	RFS4000

Type to search in tables Row Count: 15

FIGURE 361 Panic Snapshots screen

4. The screen expands to display the following parameters for each reported panic snapshot.:

Device	Displays the factory encoded MAC address assigned to the device. This is a identifier used as permanent device hardware address and cannot be modified by the controller
System Name	Lists the name assigned to each listed managed device.
Type	Displays the Brocade Mobility model of each device providing a panic to the controller.

Crash Files

Use the Crash Files screen to review files created when an access point or controller encounters a critical error or malfunction. Use crash files to troubleshoot issues specific to the device on which a crash event was generated. These are issues impacting the core (distribution layer). Once reviewed, files can be deleted or transferred for archive. Crash files can be sent to a support team to expedite issues with the reporting device.

Advanced Diagnostics

Refer to the controller's Advanced UI Diagnostics facilities to review and troubleshoot any potential issue with the controller's resident *User Interface* (UI). The UI Diagnostics screen provides a large number of diagnostic tools to enable you to effectively identify and correct issues with the controller UI. Diagnostics can also be performed at the device level for the Access Point radios and connected clients managed by the controller.

UI Debugging

Advanced Diagnostics

Use the UI Debugging screen to view debugging information for a selected device.

To review device debugging information:

1. Select **Diagnostics > Advanced** to display the UI Debugging menu options.

Once a target device has been selected its debugging information displays within the **NETCONF Viewer** by default.

The screenshot shows the 'UI Debugging' window with the 'NETCONF Viewer' tab selected. The main area displays a table of 'Real-Time NETCONF Messages'. The table has four columns: ID, Type, Operation, and Time (ms). The messages are as follows:

ID	Type	Operation	Time (ms)
667	rpc	get-config	
667	rpc-reply	data	0
668	rpc	get-config	
668	rpc-reply	data	0
669	rpc	edit-config	
669	rpc-reply	ok	0
670	rpc	get-config	
670	rpc-reply	data	0
671	rpc	get-config	
671	rpc-reply	data	0
672	rpc	get-config	
672	rpc-reply	data	0
673	rpc	edit-config	
673	rpc-reply	ok	0
674	rpc	get-config	
674	rpc-reply	data	0

At the bottom of the window, there is a search bar with the text 'Find:' and a 'Next' button. To the right, there is a 'Size:' dropdown menu set to '50', and buttons for 'Clear', 'Clipboard', and 'Clipboard All'.

FIGURE 363 UI Debugging screen - NETCONF Viewer

2. Use the **NETCONF Viewer** to review NETCONF information. NETCONF is a tag-based configuration protocol for Brocade Mobility wireless controllers. Messages are exchanged between the UI and the wireless controller using XML tags.

3. The **Real Time NETCONF Messages** area lists an XML representation of any message generated by the system. The main display area of the screen is updated in real time.
4. Refer to the **Request Response** and **Time Taken** fields on the bottom of the screen to assess the time taken by the controller to receive and respond to requests. The time is displayed in microseconds.
5. Use the **Clear** button to clear the contents of the Real Time NETCONF Messages area. Use the **Find** parameter and the **Next** button to search for message variables in the Real Time NETCONF Messages area.
6. Select **Schema Browser** to view configuration, statistics and an actions repository for a selected device.

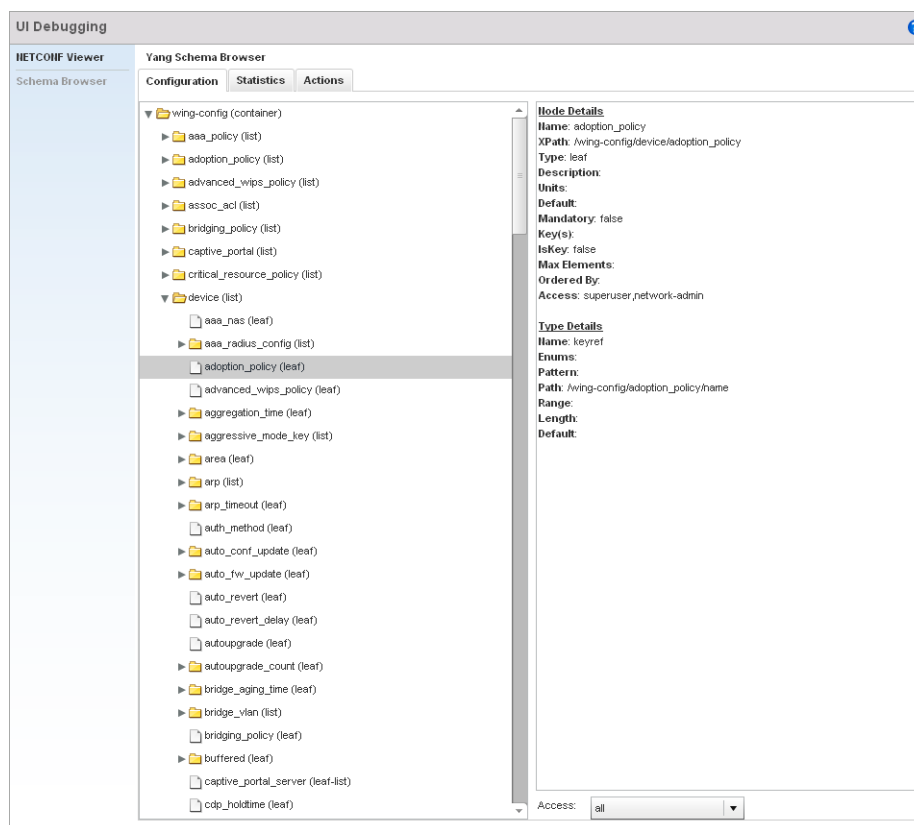


FIGURE 364 UI Debugging screen - Schema Browser Configuration tab

7. The Schema Browser is arranged into two panes (regardless of the Configuration, Statistics or Actions tab selected). The left pane allows you to navigate the schema. Selecting a node on the left pane displays the node information on the right pane. The Schema Browser does not display information in real time. It only displays the data format used on the device when last updated.
 - a. The Scheme Browser displays the **Configuration** tab by default. Expand a specific configuration parameter to review the configuration settings defined for that device parameter. The **Configuration** tab provides an ideal place to verify if device configurations differ from default settings or have been erroneously changed in respect to the device's intended configuration profile.

- b. Select the **Statistics** tab to assess performance data and statistics for a target device.

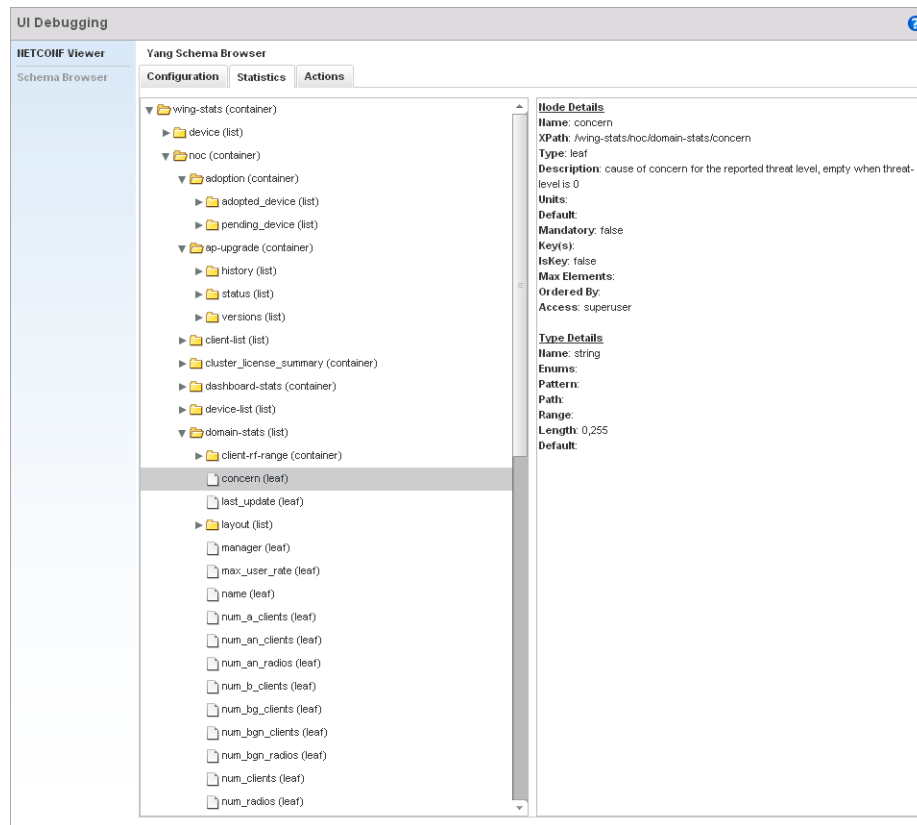


FIGURE 365 UI Debugging screen - Schema Browser Statistics tab

Use the Statistics data to assess whether the device is optimally configured in respect to its intended deployment objective. Often the roles of radio supported devices and wireless clients change as additional devices and radios are added to the managed network. Navigate amongst a target device's statistical variables to assess whether the device should be managed by a different controller profile or defined a unique configuration different from the one currently defined.

- c. Select the **Actions** tab to display schema for any action that can be configured based on an event.

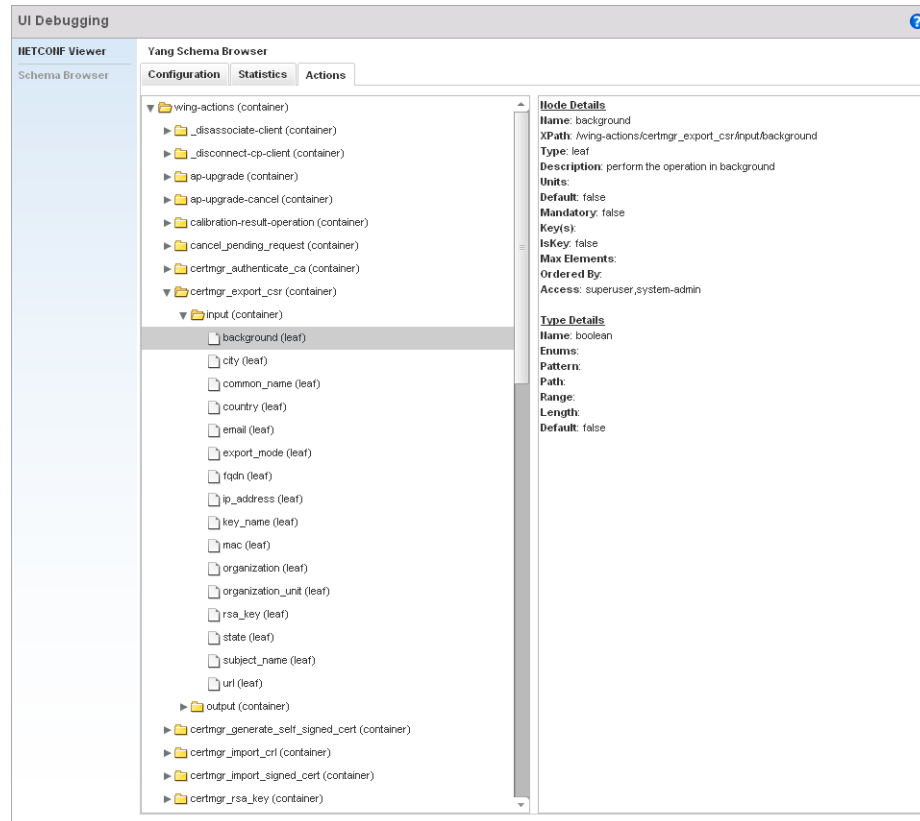


FIGURE 366 UI Debugging screen - Schema Browser Actions tab

- The left pane displays a hierarchical tree of the different actions available to the selected device. When a node is selected, its information is displayed within the right pane.

Operations

In this chapter

- [Device Operations](#) 541
- [Certificates](#) 551
- [Smart RF](#) 566

The functions within the controller's *Operations* menu allow firmware and configuration files management and certificate generation for managed devices. In a clustered environment, these operations can be performed on one controller, then propagated to each member of the cluster and onwards to the devices managed by each cluster member.

A controller certificate links identity information with a public key enclosed in the certificate. Device certificates can be imported and exported to and from the controller to a secure remote location for archive and retrieval as they are required for application to other managed devices.

Self Monitoring At Run Time RF Management (Smart RF) is a Brocade Mobility innovation designed to simplify RF configurations for new deployments, while (over time) providing on-going deployment optimization and radio performance improvements. The Smart RF functionality scans the managed network to determine the best channel and transmit power for each managed Access Point radio. Smart RF policies can be applied to specific RF Domains, to add site specific deployment configurations and self recovery values to groups of devices within pre-defined physical RF coverage areas.

Device Operations

Brocade Mobility periodically releases updated device firmware and configuration files to the Support Web site. If an Access Point's (or its associated device's) firmware is older than the version on the Web site, Brocade Mobility recommends updating to the latest firmware version for full feature functionality and optimal controller utilization. Additionally, selected devices can either have a primary or secondary firmware image applied or fallback to a selected firmware image if an error occurs in the update process.

Device update activities include:

- [Managing Firmware and Config Files](#)
- [Managing File Transfers](#)
- [Using the Controller File Browser](#)

These tasks can be performed on individual wireless controllers, Access Points and wireless clients.

Managing Firmware and Config Files

[Device Operations](#)

1. The **Device Details** screen displays by default when the **Operations** is selected from the controller's main menu bar.
2. The Device Details screen displays firmware information for a specific device selected from either the RF Domain or Network tabs on the left-hand side of the screen.

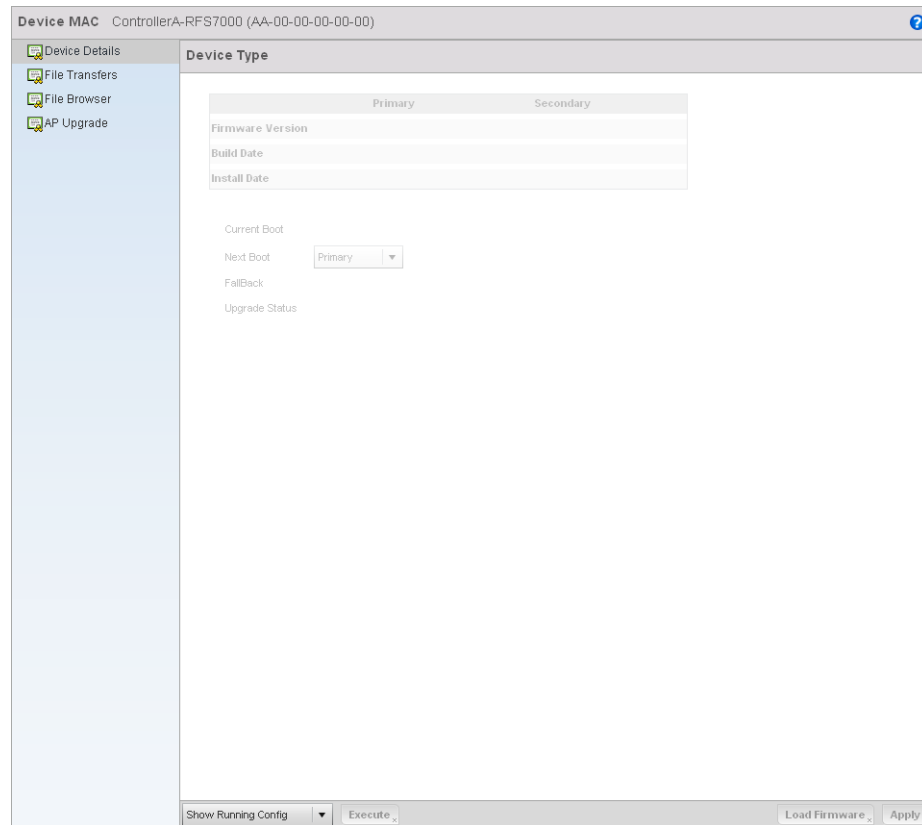


FIGURE 367 Device Details screen

3. Refer to the following to determine whether a firmware image needs to be updated for the selected device, or a device requires a restart or revert to factory default settings.

Device MAC	Displays the factory assigned hardware MAC address (in the banner of the screen) for the selected device. The Device Type also displays in the banner of the screen.
Firmware Version	Displays the primary and secondary firmware image version from the wireless controller.
Build Date	Displays the date the primary and secondary firmware image was built for the selected device.
Install Date	Displays the date the firmware was installed for the selected device.
Current Boot	Lists whether the primary or secondary firmware image is to be applied to the device the next time the device boots.
Next Boot	Use the drop-down menu to select the firmware image to boot the next time the device reboots. Select either the <i>Primary</i> or the <i>Secondary</i> image.

Fallback	Lists whether fallback is currently enabled for the selected device. When enabled, the device reverts back to the last successfully installed firmware image if something were to happen in its next firmware upgrade that would render the device inoperable.
Upgrade Status	Displays the status of the last firmware upgrade performed for each listed device managed by this controller. For information on upgrading device firmware, see <i>Upgrading Device Firmware on page 13-543</i> .
Show Startup Config	Select this option (from the drop-down menu on the bottom of the screen) to display the startup configuration of the selected device. The startup configuration is displayed in a separate window. Select the <i>Execute</i> button to perform the function.
Show Running Config	Select this option (from the drop-down menu on the bottom of the screen) to display the running configuration of the selected device. The running configuration is displayed in a separate window. Select the <i>Execute</i> button to perform the function.
Restart	Select this option (from the drop-down menu on the bottom of the screen) to restart the selected device. Selecting this option restarts the target device using its last saved configuration and does not apply factory defaults to the target device. Restarting a device resets all data collection values to zero. Select the <i>Execute</i> button to perform the function.
Restart (factory default)	Select this option (from the drop-down menu on the bottom of the screen) to restart the selected device and apply the device's factory default configuration. Selecting this option restarts the target device and reverts its configurable parameters to their factory default values. Consider exporting the device's current configuration to a secure location for archive before reverting the device to its default configuration. Select the <i>Execute</i> button to perform the function.
Halt	Select this option (from the drop-down menu on the bottom of the screen) to stop the selected device. Select the <i>Execute</i> button to perform the function.

4. For information on conducting a device firmware upgrade, see *Upgrading Device Firmware on page 13-543*. For information on file transfers, see *Managing File Transfers on page 13-544*.

Upgrading Device Firmware

[Managing Firmware and Config Files](#)

The controller has the ability to conduct firmware updates for managed devices.

To update the firmware of a managed device:

1. Select a device from the browser.
2. Select the **Load Firmware** button.



FIGURE 368 Firmware Update screen

3. By default, the **Firmware Upgrade** screen displays a **URL** field to enter the URL (destination location) of the target device firmware file. For example, `ftp://ftpuser:ftpuser@10.10.10.10/RFS4000-5.0.0.0-103R.img`

4. Enter the complete path to the firmware file for the target device.
5. If needed, select **Advanced** to expand the dialog to display network address information to the location of the target device firmware file. The number of additional fields that populate the screen is also dependent on the selected protocol.

FIGURE 369 Detailed Firmware Upgrade screen

6. Provide the following information to accurately define the location of the target device firmware file:

- | | |
|-------------------|---|
| Protocol | Select the protocol used for updating the device firmware. Available options include: <ul style="list-style-type: none"> • tftp • ftp • sftp • http • cf • usb1 • usb2 |
| Port | Use the spinner control or manually enter the value to define the port used by the protocol for firmware updates. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> . |
| IP Address | Enter IP address of the server used to update the firmware. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> . |
| Hostname | Provide the hostname of the server used to update the firmware. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> . |
| Path | Specify the path to the firmware file. Enter the complete relative path to the file on the server. |
| User Name | Define the user name used to access either a FTP or SFTP server. |
| Password | Specify the password for the user account to access a FTP or a SFTP server. |
7. Select **OK** to start the firmware update. Select **Abort** to terminate the firmware update. Select **Close** to close the upgrade popup. The upgrade continues in the background

Managing File Transfers

Device Operations

The controller can administrate files on managed devices. Transfer files from a device to this controller, to a remote server or from a remote server to the controller. An administrator can transfer logs, configurations and crash dumps.

To administrate files for managed devices:

1. Select the **Operations > Devices > File Transfers**

FIGURE 370 File Transfers screen

2. Set the following file management source and target directions as well as the configuration parameters of the required file management activity:

- Source** Select the source of the file transfer.
 Select *Server* to indicate the source of the file is a remote server.
 Select *Wireless Controller* to indicate the source of the file is the controller.
- File** If the source is *Wireless Controller*, enter the name of the file to be transferred.
- Protocol** Select the protocol for file management. Available options include:
- tftp
 - ftp
 - sftp
 - http
 - cf
 - usb1
 - usb2
- This parameter is required only when *Server* is selected as the **Source**.
- Port** Specify the port for transferring files. This option is not available for *cf*, *usb1*, and *usb2*. Enter the port number directly or use the spinner control.
 This parameter is required only when *Server* is selected as the **Source**.

- | | |
|--------------------|--|
| IP Address | Specify the IP address of the server used to transfer files. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> . If IP address of the server is provided, a Hostname is not required.
This parameter is required only when <i>Server</i> is selected as the Source . |
| Host | If needed, specify a Hostname of the server transferring the file. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> . If a hostname is provided, an IP Address is not needed.
This field is only available when <i>Server</i> is selected in the From field. |
| Path / File | Define the path to the file on the server. Enter the complete relative path to the file.
This parameter is required only when <i>Server</i> is selected as the Source . |
| User Name | Provide a user name to access a FTP or a SFTP server.
This parameter is required only when <i>Server</i> is selected as the Source , and the selected protocol is <i>ftp</i> or <i>sftp</i> . |
| Password | Provide a password to access the FTP or SFTP server.
This parameter is required only when <i>Server</i> is selected as the Source , and the selected protocol is <i>ftp</i> or <i>sftp</i> . |
| Target | Select the target destination to transfer the file. <ul style="list-style-type: none"> • Select <i>Server</i> if the destination is a remote server, then provide a URL to the location of the server resource or select Advanced and provide the same network address information described above. • Select <i>Wireless Controller</i> if the destination is the controller. |
3. Select **Copy** to begin the file transfer. Selecting **Reset** reverts the screen to its last saved configuration.

Using the Controller File Browser

Device Operations

The controller maintains a File Browser allowing an administrator to review the files residing on a controller's internal or external memory resource. Directories can be created and maintained for each File Browser location and folders and files can be moved and deleted as an administrator interprets necessary.

Keep in mind, USB1 is available on RFS4000, RFS6000 and RFS7000 model controllers, while USB2 and *Compact Flash* (CF) are only available on RFS7000 model controllers.

To administrate files for managed devices and memory resources:

1. Select the **Operations > Devices > File Browser**.

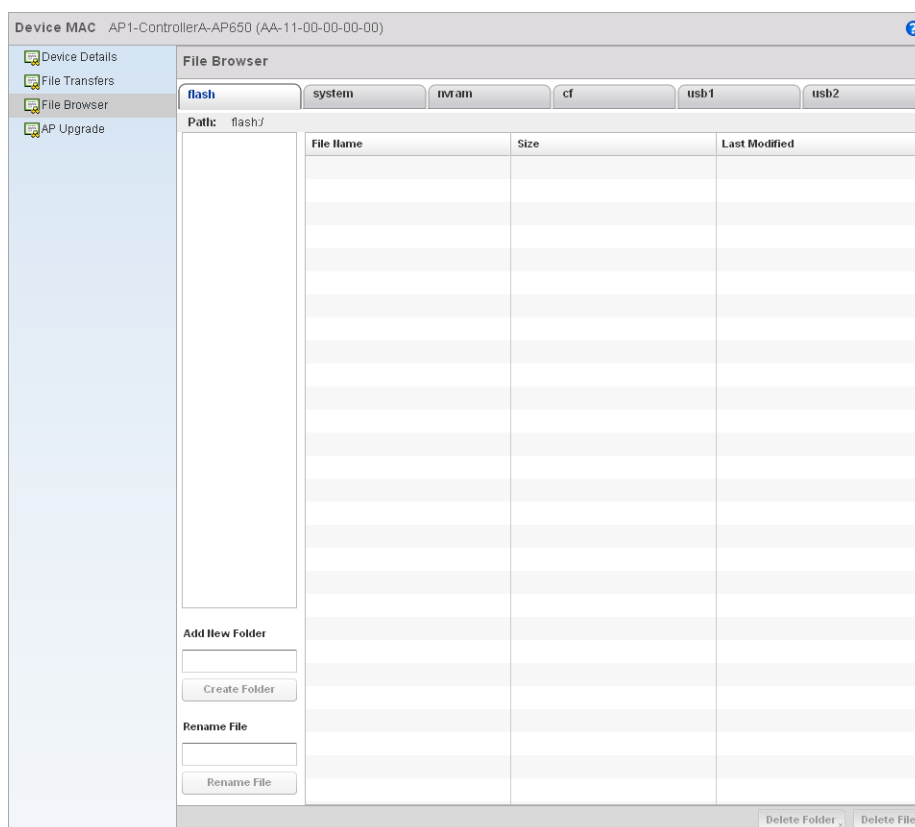


FIGURE 371 File Browser screen - flash

2. Refer to the following to determine whether a file needs to be deleted or included in a new folder for the selected internal (flash, system, nvram) or external (cf, USB1, USB2) controller memory resource. The following display for each available controller memory resource:

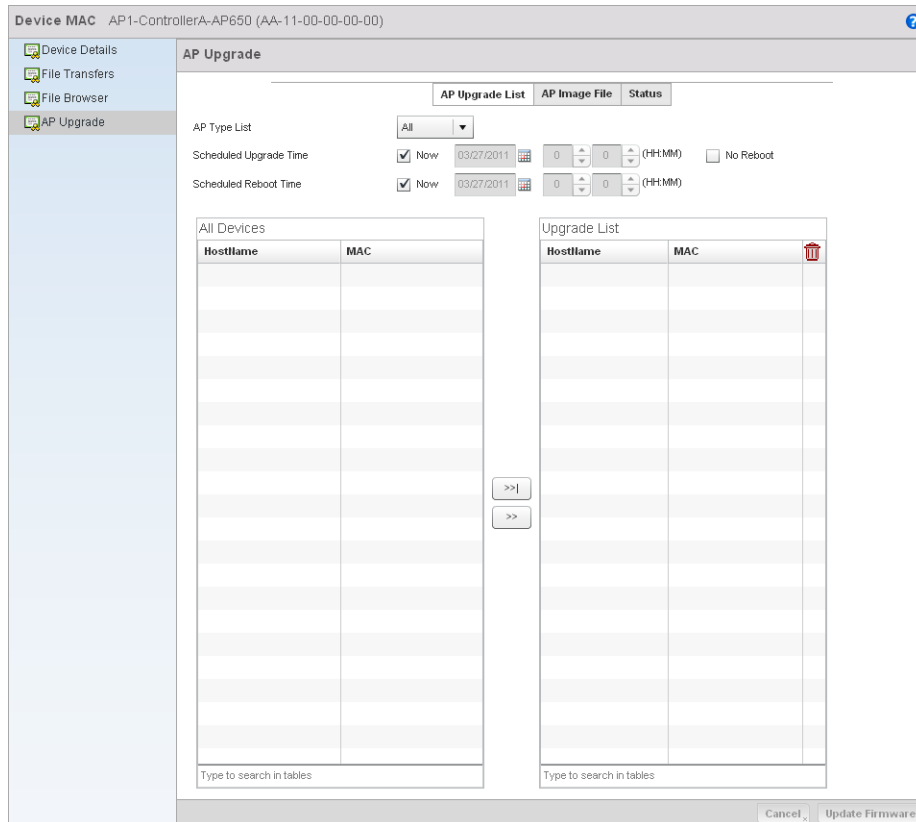
File Name	Displays the name of the file residing on the selected <i>flash</i> , <i>system</i> , <i>nvram</i> , <i>usb1</i> or <i>usb2</i> location. The name cannot be modified from this location.
Size	Displays the size of the file in kb. Use this information to help determine whether the file should be moved or deleted.
Last Modified	Lists a timestamp for the last time each listed file was modified. Use this information to determine the file's relevance or whether it should be deleted.

3. If needed, use the **Add New Folder** utility to create a folder that serves as a directory for some or all of the files for a selected controller memory resource. Once defined, select the **Create Folder** button to implement.
4. Optionally, use the **Delete Folder** or **Delete File** buttons to remove a folder or file from within the current controller memory resource.

Using the AP Upgrade Browser

[Device Operations](#)

To manage AP Upgrade configuration:

1. Select the **Operations > Devices > AP Upgrade****FIGURE 372** AP Upgrade screen

2. Select the Access Point model from the **AP Type List** drop-down to specify which model types should be available to upgrade.
3. Refer to the **Scheduled Upgrade Time** option to schedule the when the upgrade should take place. To perform an upgrade immediately, select **Now**. To schedule the upgrade to take place at a specified time, enter a date and time in the appropriate boxes.
4. Refer to the **Scheduled Reboot Time** option to schedule the when the AP should reboot. To reboot the upgraded APs immediately, select **Now**. To schedule the reboot to take place at a future time, enter a date and time in the appropriate boxes. If you do not wish for the APs to reboot after they have been upgraded, select the **No Reboot** option
5. The **All Devices** table lists available APs that match the AP Type. For each available AP, the hostname and the primary MAC Address are listed in the table. The **Upgrade List** table displays the APs selected for upgrade. For each AP, the hostname and the primary MAC Address are listed. Using the **>>|** button moves all APs listed in the All Devices table to the Upgrade List table.
6. Select the **AP Image File** tab.

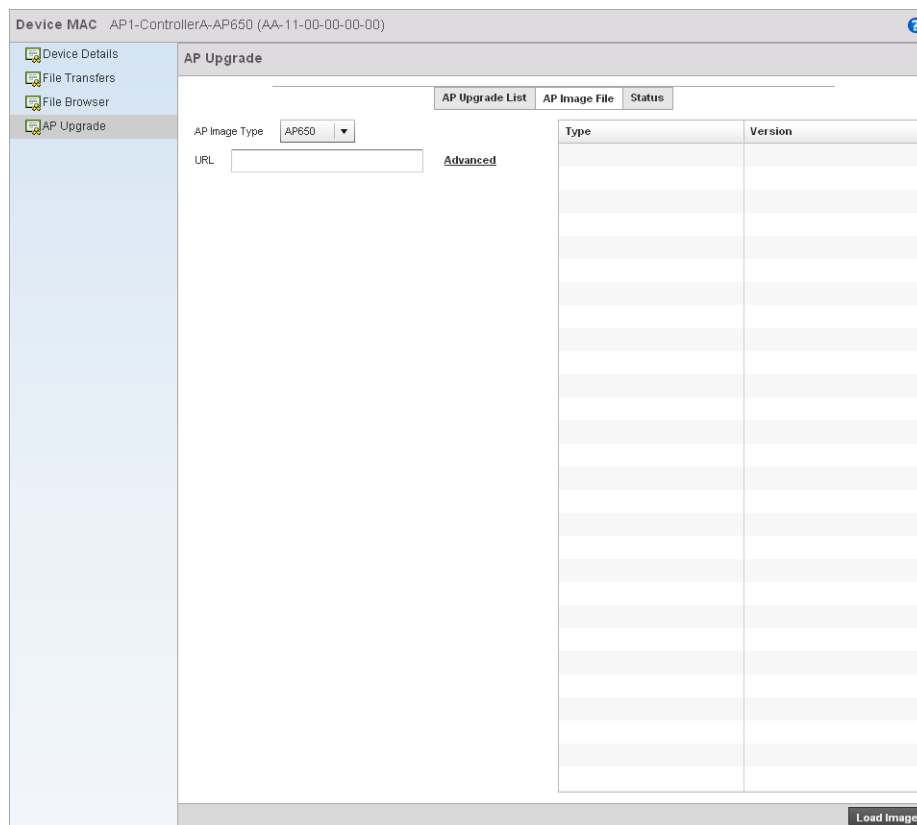


FIGURE 373 AP Upgrade screen

7. Select an Access Point model from the **AP Image Type** drop-down menu to specify which AP image types should be available during an upgrade.
8. Enter a **URL** pointing to the location of available AP image files.
9. Selecting **Advanced** will list additional options for AP image file location including protocol, host and path to the image files.

Protocol

Select the protocol for file management. Available options include:

- tftp
- ftp
- sftp
- http

Port

Specify the port for transferring files. Enter the port number directly or use the spinner control.

IP Address

Specify the IP address of the server used to transfer files. If IP address of the server is provided, a **Hostname** is not required.

Host

If needed, specify a Hostname of the server transferring the file. If a hostname is provided, an **IP Address** is not needed.

Path / File

Define the path to the file on the server. Enter the complete relative path to the file.

10. Select the **Status** tab.

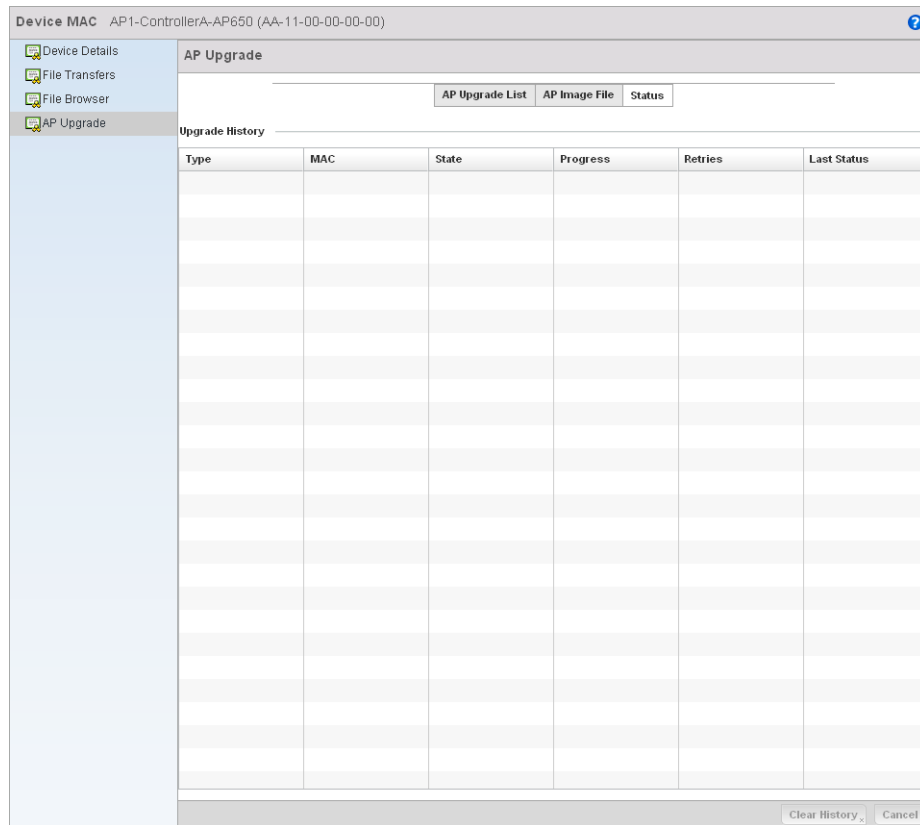


FIGURE 374 AP Upgrade screen

11. Refer to the following **Upgrade History** status information:

Type	Displays the Access Point model for each known Access Point.
MAC	Displays the primary Media Access Control (MAC) or hardware address for each known Access Point.
State	Displays the current upgrade status of each known Access Point. Possible states include: <ul style="list-style-type: none"> • Waiting • Downloading • Updating Scheduled • Reboot • Rebooting Done • Cancelled • Done • No Reboot
Progress	Displays the time of the last status update for each known Access Point undergoing an upgrade.
Retries	Displays the number of retries, if any, needed for the upgrade.
Last Status	Displays the last status update for Access Points no longer upgrading.

12. Selecting the **Clear History** button clears the current history log page for all Access Points.

13. Clicking the **Cancel** button will cancel the upgrade process for any selected Access Points that are upgrading.

Certificates

A controller certificate links identity information with a public key enclosed in the certificate.

A *certificate authority* (CA) is a network authority that issues and manages security credentials and public keys for message encryption. The CA signs all digital certificates it issues with its own private key. The corresponding public key is contained within the certificate and is called a CA certificate. A browser must contain this CA certificate in its Trusted Root Library so it can trust certificates *signed* by the CA's private key.

Depending on the public key infrastructure, the digital certificate includes the owner's public key, the certificate expiration date, the owner's name and other public key owner information.

Each certificate is digitally signed by a *trustpoint*. The trustpoint signing the certificate can be a certificate authority, corporation or individual. A trustpoint represents a CA/identity pair containing the identity of the CA, CA-specific configuration parameters and an association with an enrolled identity certificate.

SSH keys are a pair of cryptographic keys used to authenticate users instead of, or in addition to, a username/password. One key is private and the other is public key. *Secure Shell* (SSH) public key authentication can be used by a client to access managed resources, if properly configured. A RSA key pair must be generated on the client. The public portion of the key pair resides with the controller, while the private portion remains on a secure local area of the client.

For more information on the certification activities support by the controller, refer to the following:

- [Certificate Management](#)
- [RSA Key Management](#)
- [Certificate Creation](#)
- [Generating a Certificate Signing Request](#)

Certificate Management

[Certificates](#)

If not wanting to use an existing certificate or key with a selected device, an existing *stored* certificate can be leveraged from a different managed device for use with the target device. Device certificates can be imported and exported to and from the controller to a secure remote location for archive and retrieval as they are required for application to other managed devices.

To configure trustpoints for use with certificates:

1. Select **Operations > Certificates**.
2. Select a device from amongst those displayed in either the RF Domain or Network panes on the left-hand side of the screen.

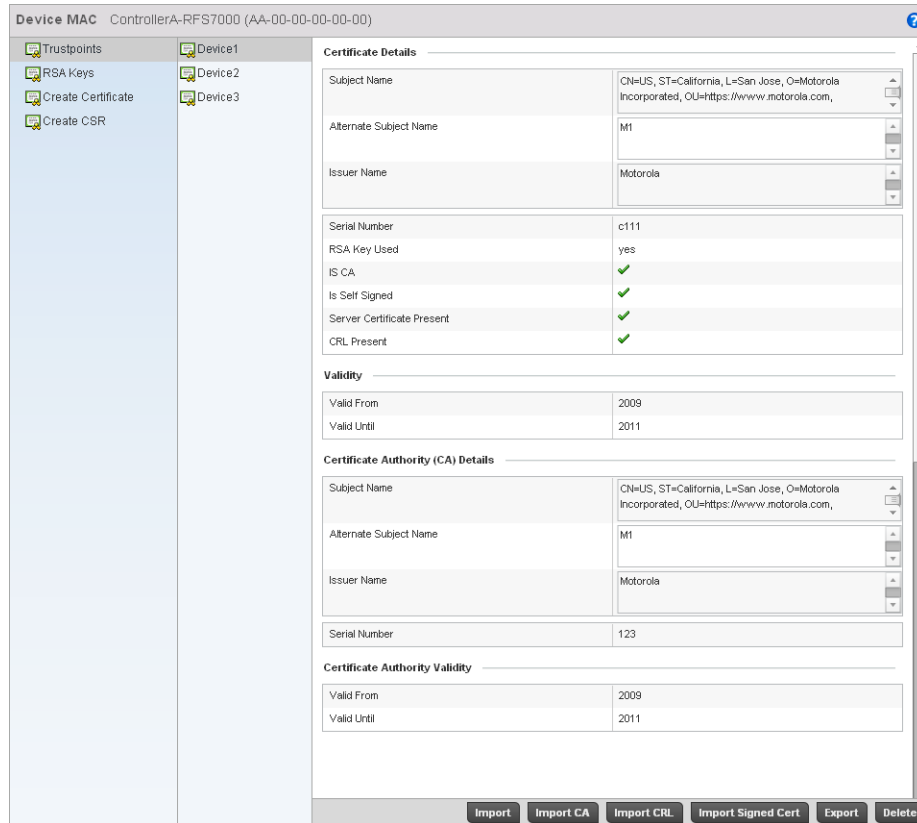


FIGURE 375 Trustpoints screen

3. The **Trustpoints** screen displays for the selected MAC address.
4. Select a device from amongst those displayed to review its certificate information.
5. Refer to the **Certificate Details** to review the certificate's properties, self-signed credentials, validity period and CA information.
6. To optionally import a certificate to the controller, select the **Import** button from the Trustpoints screen.

FIGURE 376 Import New Trustpoint screen

7. Define the following configuration parameters required for the **Import** of the trustpoint.

Trustpoint Name Enter the 32 character maximum name assigned to the target trustpoint. The trustpoint signing the certificate can be a certificate authority, corporation or individual.

Key Passphrase Define the key used by both the controller and the server (or repository) of the target trustpoint. Select the *Show* textbox to expose the characters used in the key. Leaving the *Show* checkbox unselected displays the passphrase as a series of asterisks “*”.

URL Provide the complete URL to the location of the trustpoint. If needed, select *Advanced* to expand the dialog to display network address information to the location of the target trustpoint. The number of additional fields that populate the screen is also dependent on the selected protocol.

Protocol Select the protocol used for importing the target trustpoint. Available options include:

- tftp
- ftp
- sftp
- http
- cf
- usb1
- usb2

Port Use the spinner control to set the port. This option is not valid for *cf*, *usb1* and *usb2*.

IP Address Enter IP address of the server used to import the trustpoint. This option is not valid for *cf*, *usb1* and *usb2*.

Host Provide the hostname of the server used to import the trustpoint. This option is not valid for *cf*, *usb1* and *usb2*.

Path / File Specify the path to the trustpoint. Enter the complete relative path to the file on the server.

8. Select **OK** to import the defined trustpoint. Select **Cancel** to revert the screen to its last saved configuration.

9. To optionally import a CA certificate to the controller, select the **Import CA** button from the Trustpoints screen.

A *certificate authority (CA)* is a network authority that issues and manages security credentials and public keys for message encryption. The CA signs all digital certificates it issues with its own private key. The corresponding public key is contained within the certificate and is called a *CA certificate*.

FIGURE 377 Import CA Certificate screen

10. Define the following configuration parameters required for the **Import** of the CA certificate:

Trustpoint Name	Enter the 32 character maximum name assigned to the target trustpoint signing the certificate. A trustpoint represents a CA/identity pair containing the identity of the CA, CA-specific configuration parameters and an association with an enrolled identity certificate.
From Network	Select the <i>From Network</i> radio button to provide network address information to the location of the target CA certificate. The number of additional fields populating the screen is dependent on the selected protocol.
Cut and Paste	Select the <i>Cut and Paste</i> radio button to simply copy an existing CA certificate into the cut and past field. When pasting a valid CA certificate, no additional network address information is required.
Protocol	Select the protocol used for importing the target CA certificate. Available options include: <ul style="list-style-type: none"> • tftp • ftp • sftp • http • cf • usb1 • usb2
Port	Use the spinner control to set the port. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> .
IP Address	Enter IP address of the server used to import the CA certificate. This option is not valid for <i>cf</i> , <i>usb1</i> , and <i>usb2</i> .
Hostname	Provide the hostname of the server used to import the CA certificate. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Path	Specify the path to the CA certificate. Enter the complete relative path to the file on the server.

11. Select **OK** to import the defined CA certificate. Select **Cancel** to revert the screen to its last saved configuration.

12. To optionally import a CRL to the controller, select the **Import CRL** button from the Trustpoints screen.

If a certificate displays within the Certificate Management screen with a CRL, that CRL can be imported into the controller. A *certificate revocation list* (CRL) is a list of certificates that have been revoked or are no longer valid. A certificate can be revoked if the *certificate authority* (CA) had improperly issued a certificate, or if a private-key is compromised. The most common reason for revocation is the user no longer being in sole possession of the private key.

FIGURE 378 Import CRL screen

13. Define the following configuration parameters required for the **Import** of the CRL:

Trustpoint Name	Enter the 32 character maximum name assigned to the target trustpoint signing the certificate. A trustpoint represents a CA/identity pair containing the identity of the CA, CA-specific configuration parameters, and an association with an enrolled identity certificate.
From Network	Select the <i>From Network</i> radio button to provide network address information to the location of the target CRL. The number of additional fields that populate the screen is also dependent on the selected protocol. This is the default setting.
Cut and Paste	Select the <i>Cut and Paste</i> radio button to simply copy an existing CRL into the cut and past field. When pasting a CRL, no additional network address information is required.
URL	Provide the complete URL to the location of the CRL. If needed, select <i>Advanced</i> to expand the dialog to display network address information to the location of the target CRL. The number of additional fields that populate the screen is also dependent on the selected protocol.
Protocol	Select the protocol used for importing the CRL. Available options include: <ul style="list-style-type: none"> • tftp • ftp • sftp • http • cf • usb1 • usb2

Port	Use the spinner control to set the port. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
IP Address	Enter IP address of the server used to import the CRL. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Hostname	Provide the hostname of the server used to import the CRL. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Path	Specify the path to the CRL. Enter the complete relative path to the file on the server.

14. Select **OK** to import the CRL. Select **Cancel** to revert the screen to its last saved configuration.
15. To import a signed certificate to the controller, select the **Import Signed Cert** button from the Trustpoints screen.

Signed certificates (or root certificates) avoid the use of public or private CAs. A self signed certificate is an identity certificate signed by its own creator. Thus, the certificate creator also signs off on its legitimacy. The lack of mistakes or corruption in the issuance of self signed certificates is central.

Self-signed certificates cannot be revoked, which may allow an attacker who has already gained controller access to monitor and inject data into a connection to spoof an identity if a private key has been compromised. However, CAs have the ability to revoke a compromised certificate, which prevents its further use.

FIGURE 379 Import Signed Cert screen

16. Define the following configuration parameters required for the **Import** of the CA certificate:

Certificate Name	Enter the 32 character maximum name of the trustpoint with which the certificate should be associated
From Network	Select the <i>From Network</i> radio button to provide network address information to the location of the target signed certificate. The number of additional fields that populate the screen is also dependent on the selected protocol. This is the default setting.
Cut and Paste	Select the <i>Cut and Paste</i> radio button to simply copy an existing signed certificate into the cut and past field. When pasting a signed certificate, no additional network address information is required.
URL	Provide the complete URL to the location of the signed certificate. If needed, select <i>Advanced</i> to expand the dialog to display network address information to the location of the signed certificate. The number of additional fields that populate the screen is dependent on the selected protocol.
Protocol	Select the protocol used for importing the target signed certificate. Available options include: <ul style="list-style-type: none"> • tftp • ftp • sftp • http • cf • usb1 • usb2
Port	Use the spinner control to set the port. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
IP Address	Enter IP address of the server used to import the signed certificate. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Hostname	Provide the hostname of the server used to import the signed certificate. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Path	Specify the path to the signed certificate. Enter the complete relative path to the file on the server.

17. Select **OK** to import the signed certificate. Select **Cancel** to revert the screen to its last saved configuration

18. To optionally export a trustpoint from the controller to a remote location, select the *Export* button from the Trustpoints screen.

Once a certificate has been generated on the controller's authentication server, export the self signed certificate. A digital CA certificate is different from a self signed certificate. The CA certificate contains the public and private key pairs. The self certificate only contains a public key. Export the self certificate for publication on a Web server or file server for certificate deployment or export it in to an Active Directory Group Policy for automatic root certificate deployment.

Additionally export the key to a redundant RADIUS server so it can be imported without generating a second key. If there's more than one RADIUS authentication server, export the certificate and don't generate a second key unless you want to deploy two root certificates.

FIGURE 380 Export Trustpoint screen

19. Define the following configuration parameters required for the *Export* of the trustpoint.

Trustpoint Name	Enter the 32 character maximum name assigned to the target trustpoint. The trustpoint signing the certificate can be a certificate authority, corporation or individual.
Key Passphrase	Define the key used by both the controller and the server (or repository) of the target trustpoint. Select the Show textbox to expose the actual characters used in the key. Leaving the Show checkbox unselected displays the passphrase as a series of asterisks “*”.
URL	Provide the complete URL to the location of the trustpoint. If needed, select <i>Advanced</i> to expand the dialog to display network address information to the location of the target trustpoint. The number of additional fields that populate the screen is also dependent on the selected protocol.
Protocol	Select the protocol used for exporting the target trustpoint. Available options include: <ul style="list-style-type: none"> • tftp • ftp • sftp • http • cf • usb1 • usb2
Port	Use the spinner control to set the port. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
IP Address	Enter IP address of the server used to export the trustpoint. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Hostname	Provide the hostname of the server used to export the trustpoint. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Path	Specify the path to the trustpoint. Enter the complete relative path to the file on the server.

20. Select **OK** to export the defined trustpoint. Select **Cancel** to revert the screen to its last saved configuration.

- To optionally delete a trustpoint, select the **Delete** button from the Trustpoints screen. Provide the trustpoint name within the **Delete Trustpoint** screen and optionally select the **Delete RSA Key** checkbox to remove the RSA key along with the trustpoint. Select **OK** to proceed with the deletion, or **Cancel** to revert to the last saved configuration.

RSA Key Management

Certificates

Refer to the RSA Keys screen to review existing RSA key configurations applied to managed devices. If an existing key does not meet the needs of a pending certificate request, generate a new key or import or export an existing key to and from a remote location.

Rivest, Shamir, and Adleman (RSA) is an algorithm for public key cryptography. It's an algorithm used for certificate signing and encryption. When a device trustpoint is created, the RSA key is the private key used with the trustpoint.

To review existing device RSA key configurations, generate additional keys or import/export keys to and from remote locations:

- Select **Operations > Certificates**.
- Select a device from amongst those displayed in either the RF Domain or Network panes on the left-hand side of the screen.
- Select **RSA Keys**.

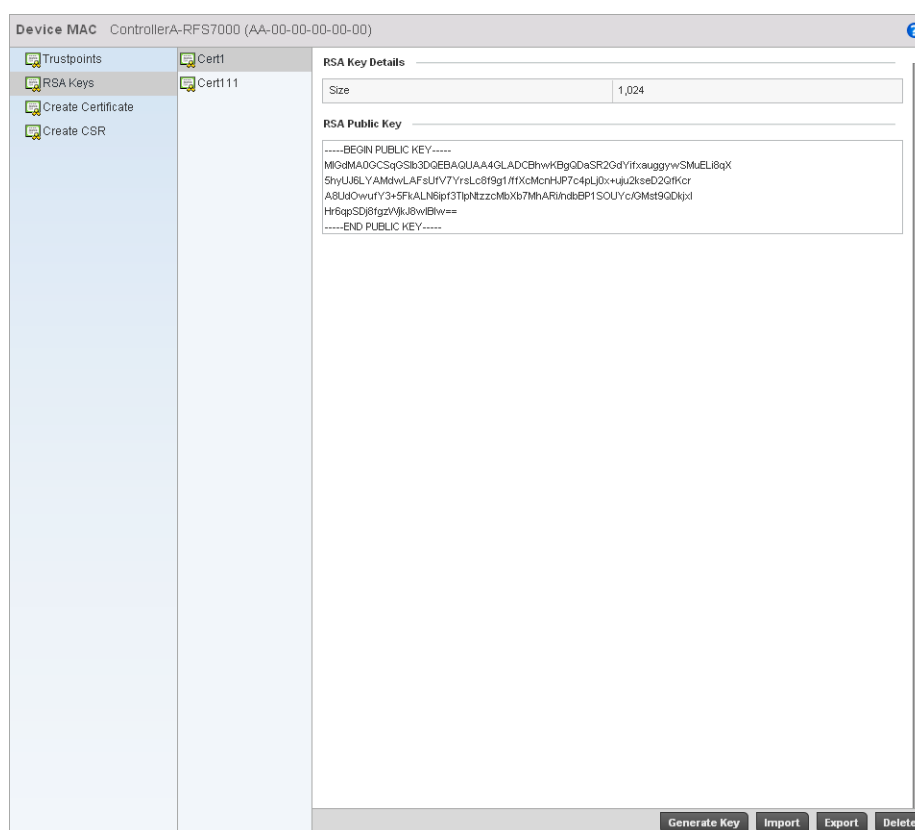


FIGURE 381 RSA Keys screen

Each key can have its size and character syntax displayed. Once reviewed, optionally generate a new RSA key, import a key from a selected device, export a key from the controller to a remote location or delete a key from a selected device.

4. Select **Generate Key** to create a new key with a defined size.

FIGURE 382 Generate RSA Key screen

5. Select **OK** to generate the RSA key. Select **Cancel** to revert the screen to its last saved configuration.

Key Name Enter the 32 character maximum name assigned to the RSA key.

Key Size Use the spinner control to set the size of the key (between 1,024 - 2,048 bits). Brocade Mobility recommends leaving this value at the default setting of 1024 to ensure optimum functionality.

6. To optionally import a CA certificate to the controller, select the **Import** button from the RSA Keys screen.

FIGURE 383 Import New RSA Key screen

7. Define the following configuration parameters required for the Import of the RSA key:

Key Name	Enter the 32 character maximum name assigned to identify the RSA key.
Key Passphrase	Define the key used by both the controller and the server (or repository) of the target RSA key. Select the <i>Show</i> textbox to expose the actual characters used in the passphrase. Leaving the Show checkbox unselected displays the passphrase as a series of asterisks “*”.
URL	Provide the complete URL to the location of the RSA key. If needed, select <i>Advanced</i> to expand the dialog to display network address information to the location of the target key. The number of additional fields that populate the screen is also dependent on the selected protocol.
Protocol	Select the protocol used for importing the target key. Available options include: <ul style="list-style-type: none"> • tftp • ftp • sftp • http • cf • usb1 • usb2
Port	Use the spinner control to set the port. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
IP Address	Enter IP address of the server used to import the RSA key. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Hostname	Provide the hostname of the server used to import the RSA key. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Path	Specify the path to the RSA key. Enter the complete relative path to the key on the server.

8. Select **OK** to import the defined RSA key. Select **Cancel** to revert the screen to its last saved configuration.
9. To optionally export a RSA key from the controller to a remote location, select the *Export* button from the RSA Keys screen.
10. Export the key to a redundant RADIUS server so it can be imported without generating a second key. If there's more than one RADIUS authentication server, export the certificate and don't generate a second key unless you want to deploy two root certificates.

FIGURE 384 Export RSA Key screen

11. Define the following configuration parameters required for the Export of the RSA key.

Key Name	Enter the 32 character maximum name assigned to the RSA key.
Key Passphrase	Define the key passphrase used by both the controller and the server. Select the Show textbox to expose the actual characters used in the passphrase. Leaving the Show checkbox unselected displays the passphrase as a series of asterisks “*”.
URL	Provide the complete URL to the location of the key. If needed, select <i>Advanced</i> to expand the dialog to display network address information to the location of the target key. The number of additional fields that populate the screen is also dependent on the selected protocol.
Protocol	Select the protocol used for exporting the RSA key. Available options include: <ul style="list-style-type: none"> • tftp • ftp • sftp • http • cf • usb1 • usb2
Port	Use the spinner control to set the port. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
IP Address	Enter IP address of the server used to export the RSA key. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Hostname	Provide the hostname of the server used to export the RSA key. This option is not valid for <i>cf</i> , <i>usb1</i> and <i>usb2</i> .
Path	Specify the path to the key. Enter the complete relative path to the key on the server.

12. Select **OK** to export the defined RSA key. Select **Cancel** to revert the screen to the last saved configuration.

13. To optionally delete a key, select the Delete button from within the RSA Keys screen. Provide the key name within the Delete RSA Key screen and select the **Delete Certificates** checkbox to remove the certificate the key supported. Select **OK** to proceed with the deletion, or **Cancel** to revert to the last saved configuration.

Certificate Creation

Certificates

The Certificate Management screen provides the facility for creating new self-signed certificates. Self signed certificates (often referred to as root certificates) do not use public or private CAs. A self signed certificate is a certificate signed by its own creator, with the certificate creator responsible for its legitimacy.

To create a self-signed certificate that can be applied to a managed device:

1. Select **Operations > Certificates**.
2. Select a device from amongst those displayed in either the RF Domain or Network panes on the left-hand side of the screen.
3. Select **Create Certificate**.

The screenshot shows a web interface for creating a self-signed certificate. The title bar reads "Device MAC ControllerA-RFS7000 (AA-00-00-00-00-00)". On the left is a navigation menu with "Trustpoints", "RSA Keys", "Create Certificate" (highlighted), and "Create CSR". The main content area is titled "Create New Self-Signed Certificate". It contains the following fields and options:

- Certificate Name**: A text input field with a red asterisk.
- RSA Key**: A radio button group with "Create New" selected and "Use Existing" unselected. Below it is a text input field with a red asterisk, a dropdown menu showing "1024", and a note "(1,024 to 2,048 bits)".
- Certificate Subject Name**: A radio button group with "auto-generate" selected and "user-configured" unselected.
- Country (C)**: Text input field.
- State (ST)**: Text input field.
- City (L)**: Text input field.
- Organization (O)**: Text input field.
- Organizational Unit (OU)**: Text input field.
- Common Name (CN)**: Text input field.
- Additional Credentials**:
 - Email Address**: Text input field.
 - Domain Name**: Text input field.
 - IP Address**: Text input field with a dotted pattern.

A "Generate Certificate" button is located at the bottom right of the form.

FIGURE 385 Create Certificate screen

4. Define the following configuration parameters required to Create New Self-Signed Certificate:

- | | |
|--------------------------------|---|
| Certificate Name | Enter the 32 character maximum name assigned to identify the name of the trustpoint associated with the certificate. A trustpoint represents a CA/identity pair containing the identity of the CA, CA-specific configuration parameters, and an association with an enrolled identity certificate. |
| Use an Existing RSA Key | Select the radio button and use the drop-down menu to select the existing key used by both the controller and the server (or repository) of the target RSA key. |
| Create a New RSA Key | To create a new RSA key, select the radio button to define 32 character name used to identify the RSA key. Use the spinner control to set the size of the key (between 1,024 - 2,048 bits). Brocade Mobility recommends leaving this value at the default setting of 1024 to ensure optimum functionality. For more information on creating a new RSA key, see <i>RSA Key Management on page 13-559</i> . |

5. Set the following Certificate Subject Name parameters required for the creation of the certificate:

- | | |
|---------------------------------|---|
| Certificate Subject Name | Select either the <i>auto-generate</i> radio button to automatically create the certificate's subject credentials or select <i>user-defined</i> to manually enter the credentials of the self signed certificate. The default setting is auto-generate. |
| Country (C) | Define the Country used in the certificate. The field can be modified by the user to other values. This is a required field and cannot exceed 2 characters. |
| State (ST) | Enter a State/Prov. for the state or province name used in the certificate. This is a required field. |
| City (L) | Enter a City to represent the city name used in the certificate. This is a required field. |
| Organization (O) | Define an Organization for the organization used in the certificate. This is a required field. |
| Organizational Unit (OU) | Enter an Org. Unit for the name of the organization unit used in the certificate. This is a required field. |
| Common Name (CN) | If there's a common name (IP address) for the organizational unit issuing the certificate, enter it here. |

6. Select the following Additional Credentials required for the generation of the self signed certificate:

- | | |
|----------------------|---|
| Email Address | Provide an email address used as the contact address for issues relating to this certificate request. |
| Domain Name) | Enter a <i>fully qualified domain name</i> (FQDN) is an unambiguous domain name that specifies the node's position in the DNS tree hierarchy absolutely. To distinguish an FQDN from a regular domain name, a trailing period is added. ex: somehost.example.com. An FQDN differs from a regular domain name by its absoluteness; as a suffix is not added. |
| IP Address | Specify the controller IP address used as the controller destination for certificate requests. |

7. Select the **Generate Certificate** button at the bottom of the Create Certificate screen to produce the certificate.

Generating a Certificate Signing Request

[Certificates](#)

A *certificate signing request (CSR)* is a message from a requestor to a certificate authority to apply for a digital identity certificate. The CSR is composed of a block of encrypted text generated on the server the certificate will be used on. It contains information included in the certificate, including organization name, common name (domain name), locality, and country.

A RSA key must be either created or applied to the certificate request before the certificate can be generated. A private key is not included in the CSR, but is used to digitally sign the completed request. The certificate created with a particular CSR only worked with the private key generated with it. If the private key is lost, the certificate is no longer functional. The CSR can be accompanied by other identity credentials required by the certificate authority, and the certificate authority maintains the right to contact the applicant for additional information.

If the request is successful, the CA sends an identity certificate digitally signed with the private key of the CA.

To create a CSR:

1. Select **Operations > Certificates**.
2. Select a device from amongst those displayed in either the RF Domain or Network panes on the left-hand side of the screen.
3. Select **Create CSR**.

The screenshot shows the 'Create New Certificate Signing Request (CSR)' interface for a device named 'ControllerA-RFS7000 (AA-00-00-00-00-00)'. The interface is divided into several sections:

- Trustpoints:** A sidebar on the left contains links for 'Trustpoints', 'RSA Keys', 'Create Certificate', and 'Create CSR' (which is highlighted).
- Header:** 'Device MAC ControllerA-RFS7000 (AA-00-00-00-00-00)' and a help icon.
- Form Section:**
 - RSA Key:** Radio buttons for 'Create New' (selected) and 'Use Existing'. A text input field with a dropdown menu showing '1024' and '(1,024 to 2,048 bits)'.
 - Certificate Subject Name:** Radio buttons for 'auto-generate' (selected) and 'user-configured'.
 - Fields:** Text input fields for 'Country (C)', 'State (ST)', 'City (L)', 'Organization (O)', 'Organizational Unit (OU)', and 'Common Name (CN)'.
 - Additional Credentials:** Text input fields for 'Email Address', 'Domain Name', and 'IP Address'.
- Footer:** A 'Generate CSR' button.

FIGURE 386 Create CSR screen

4. Define the following configuration parameters required to **Create New Certificate Signing Request (CSR)**:

Use an Existing RSA Key Select the radio button and use the drop-down menu to select the existing key used by both the controller and the server (or repository) of the target RSA key.

Create a New RSA Key To create a new RSA key, select the radio button to define 32 character name used to identify the RSA key. Use the spinner control to set the size of the key (between 1,024 - 2,048 bits). Brocade Mobility recommends leaving this value at the default setting of 1024 to ensure optimum functionality. For more information on creating a new RSA key, see *RSA Key Management on page 13-559*.

5. Set the following Certificate Subject Name parameters required for the creation of the certificate:

Certificate Subject Name Select either the *auto-generate* radio button to automatically create the certificate's subject credentials or select *user-defined* to manually enter the credentials of the self signed certificate. The default setting is auto-generate.

Country (C) Define the Country used in the CSR. The field can be modified by the user to other values. This is a required field and must not exceed 2 characters.

State (ST) Enter a State/Prov. for the state or province name used in the CSR. This is a required field.

City (L) Enter a City to represent the city name used in the CSR. This is a required field.

Organization (O) Define an Organization for the organization used in the CSR. This is a required field.

Organizational Unit (OU) Enter an Org. Unit for the name of the organization unit used in the CSR. This is a required field.

Common Name (CN) If there's a common name (IP address) for the organizational unit issuing the certificate, enter it here.

6. Select the following **Additional Credentials** required for the generation of the CSR:

Email Address Provide an email address used as the contact address for issues relating to this CSR.

Domain Name) Enter a *fully qualified domain name* (FQDN) is an unambiguous domain name that specifies the node's position in the DNS tree hierarchy absolutely. A trailing period is added to distinguish an FQDN from a regular domain name. For example, *somehost.example.com*. An FQDN differs from a regular domain name by its absoluteness, since a suffix is not added.

IP Address Specify the controller IP address used as the controller destination for certificate requests.

7. Select the **Generate CSR** button at the bottom of the screen to produce the CSR.

Smart RF

Self Monitoring At Run Time RF Management (Smart RF) is a Brocade Mobility innovation designed to simplify RF configurations for new deployments, while (over time) providing on-going deployment optimization and radio performance improvements.

The Smart RF functionality scans the managed network to determine the best channel and transmit power for each wireless controller managed Access Point radio. Smart RF policies can be applied to specific RF Domains, to apply site specific deployment configurations and self recovery values to groups of devices within pre-defined physical RF coverage areas.

Smart RF also provides self recovery functions by monitoring the managed network in real-time and provides automatic mitigation from potentially problematic events such as radio interference, coverage holes and radio failures. Smart RF employs self recovery to enable a WLAN to better maintain wireless client performance and site coverage during dynamic RF environment changes, which typically require manual reconfiguration to resolve.

Smart RF is supported in standalone and clustered environments. In standalone environments, the individual controller manages the calibration and monitoring phases. In clustered environments, a single controller is elected a Smart Scan master and the remaining cluster members operate as Smart RF clients. In cluster operation, the Smart Scan master coordinates calibration and configuration and during the monitoring phase receives information from the Smart RF clients. Smart RF calibration can be triggered manually or continues at run-time, all the time.

Smart RF is supported on RFS4000, RFS6000 and RFS7000 model wireless controllers managing Brocade Mobility 650 Access Point, AP6511 or AP-7131 (adaptive mode) Access Points in either standalone or clustered environments.

NOTE

For AP7131 series Access Points, Smart RF should only be used with the façade antenna, and for an Brocade Mobility 650 Access Point, Smart RF should only be used with internal antenna models.

Within the Operations node, Smart RF is managed within selected RF Domains, using the Access Points that comprise the RF Domain and their respective radio and channel configurations as the basis to conduct Smart RF calibration operations.

Managing Smart RF for an RF Domain

Smart RF

When calibration is initiated, Smart RF instructs adopted radios (within a selected controller RF Domain) to beacon on a specific legal channel, using a specific transmit power setting. Smart RF measures the signal strength of each beacon received from both managed and unmanaged neighboring APs to define a RF map of the neighboring radio coverage area. Smart RF uses this information to calculate each managed radio's RF configuration as well as assign radio roles, channel and power.

Within a well planned RF Domain, any associated radio should be reachable by at least one other radio. The Smart RF feature records signals received from its neighbors. Access Point to Access Point distance is recorded in terms of signal attenuation. The information is used during channel assignment to minimize interference.

To conduct Smart RF calibration for an RF Domain:

1. Select **Operations > Smart RF**.
2. Expand the System mode in the upper, left-hand, side of the controller user interface to display the RF Domains available for Smart RF calibration.
3. Select a RF Domain from amongst those displayed.

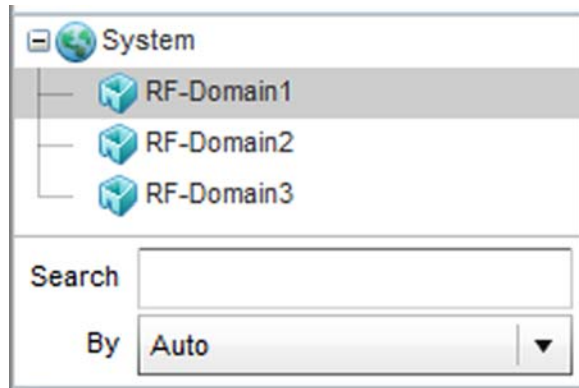


FIGURE 387 Smart RF screen - System Browser

4. The Smart RF screen displays information specific to the devices within the selected RF Domain using data from the last interactive calibration.

The screenshot shows the Smart RF screen for 'RF-Domain1'. It features a table with the following columns: AP MAC Address, MAC Address, Radio Index, Old Channel, Channel, Old Power, Power, Smart Sensor, State, and Type. The table is currently empty. Below the table, there is a warning icon and the text 'Table does not have any data to display until Calibration is run'. A search box labeled 'Type to search in tables' is present, along with a 'Row Count: 0' indicator. At the bottom, there are several buttons: Refresh, Clear Config, Clear History, Interactive Calibration, Calibration Result Actions, Run Calibration, and Stop Calibration.

FIGURE 388 Smart RF screen

5. Refer to the following to determine whether a Smart RF calibration or an interactive calibration is required:

AP MAC Address	Displays the hardware encoded MAC address assigned to each Access Point radio within the selected RF Domain. This value cannot be modified as part of a calibration activity.
MAC Address	Displays the hardware encoded MAC address assigned to each Access Point radio within the selected RF Domain. This value cannot be modified as part of a calibration activity.
Radio Index	Displays a numerical index assigned to each listed Access Point radio when it was added to the managed network. This index helps distinguish this radio from others within this RF Domain with similar configurations. This value is not subject to change as a result of a calibration activity, but each listed radio index can be used in Smart RF calibration.
Old Channel	Lists the channel originally assigned to each listed Access Point MAC address within this RF Domain. This value may have been changed as part of an Interactive Calibration process applied to this RF Domain. Compare this Old Channel against the Channel value to right of it (in the table) to determine whether a new channel assignment was warranted to compensate for a coverage hole.
Channel	Lists the current channel assignment for each listed Access Point, as potentially updated by an Interactive Calibration. Use this data to determine whether a channel assignment was modified as part of an Interactive Calibration. If a revision was made to the channel assignment, a coverage hole was detected on the channel as a result of a potentially failed or under performing Access Point radio within this RF Domain.
Old Power	Lists the transmit power assigned to each listed Access Point MAC address within this RF Domain. The power level may have been increased or decreased as part of an Interactive Calibration process applied to this RF Domain. Compare this Old Power level against the Power value to right of it (in the table) to determine whether a new power level was warranted to compensate for a coverage hole.
Power	This column displays the transmit power level for the listed Access Point MAC address after an Interactive Calibration resulted in an adjustment. This is the new power level defined by Smart RF to compensate for a coverage hole.
Smart Sensor	Defines whether a listed Access Point is smart sensor on behalf of the other Access Point radios comprising the RF Domain.
State	Displays the current state of the Smart RF managed Access Point radio. Possible states include: <i>Normal</i> , <i>Offline</i> and <i>Sensor</i> .
Type	Displays the radio type (802.11an, 802.11bgn etc.) of each listed Access Point radio within the selected RF Domain.

6. Select the **Refresh** button to (as needed) to update the contents of the Smart RF screen and the attributes of the devices within the selected RF Domain.
7. Select the **Interactive Calibration** button to initiate a Smart RF calibration using the Access Points within the selected RF Domain. The results of the calibration display within the Smart RF screen. Of particular interest are the channel and power adjustments made by the controller's Smart RF module. Expand the screen to display the Event Monitor to track the progress of the Interactive Calibration.
8. Select the **Save Calibration Result** screen to launch a sub screen used to determine the actions taken based on the results of the Interactive Calibration. The results of an Interactive calibration are not applied to radios directly, the administrator has the choice to select one of following options:

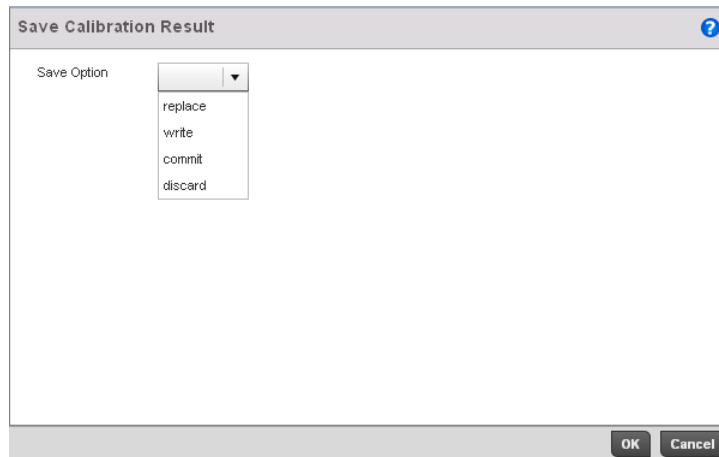


FIGURE 389 Save Calibration Result screen

Replace	Overwrites the current channel and power values with new channel power values the Interactive Calibration has calculated.
Write	Writes the new channel and power values to the radios under their respective device configurations.
Discard	Discards the results of the Interactive Calibration without applying them to their respective devices.
Commit	Commits the Smart RF module Interactive Calibration results to their respective Access Point radios.

9. Select the **Run Calibration** option to initiate a calibration. New channel and power values are applied to radios, they are not written to the running-configuration.

These values are dynamic and may keep changing during the course of the run-time monitoring and calibration the Smart RF module keeps performing to continually maintain good coverage. Unlike an Interactive Calibration, the Smart RF screen is not populated with the changes needed on Access Point radios to remedy a detected coverage hole. Expand the screen to display the Event Monitor to track the progress of the calibration.

10. The calibration process can be stopped by selecting the **Stop Calibration** button.

Statistics

In this chapter

- [System Statistics](#) 571
- [RF Domain Statistics](#) 579
- [Access Point Statistics](#) 593
- [Wireless Controller Statistics](#) 649
- [Wireless Client Statistics](#) 713

This chapter describes statistics displayed by the controller GUI. Statistics are available for both the controller and its managed devices.

A Smart RF statistical history is available to assess adjustments made to device configurations to compensate for detected coverage holes or device failures.

Access Point statistics can be exclusively displayed to validate connected Access Points, their VLAN assignments and their current authentication and encryption schemes.

Controller statistics display detailed information about controller peers, controller health, device inventories, wireless clients associations, adopted AP information, rogue APs and WLANs.

Wireless client statistics are available for an overview of client health. Wireless client statistics includes RF quality, traffic utilization and user details. Use this information to assess if configuration changes are required to improve network performance.

System Statistics

The **System** screen displays information supporting managed devices. Use this information to obtain an overall view of the state of the devices in the network. The data is organized as follows:

- [Health](#)
- [Inventory](#)
- [Adopted Devices](#)
- [Pending Adoptions](#)
- [Licenses](#)

Health

[System Statistics](#)

The *Health* screen displays the overall performance of the wireless controller managed network. This includes information on the device availability, overall RF quality, resource utilization and network threat perception.

To display the health of the wireless controller managed network:

1. Select the **Statistics** menu from the Web UI.
2. Select the **System** node from the left navigation pane.
3. Select **Health** from the left-hand side of the UI.

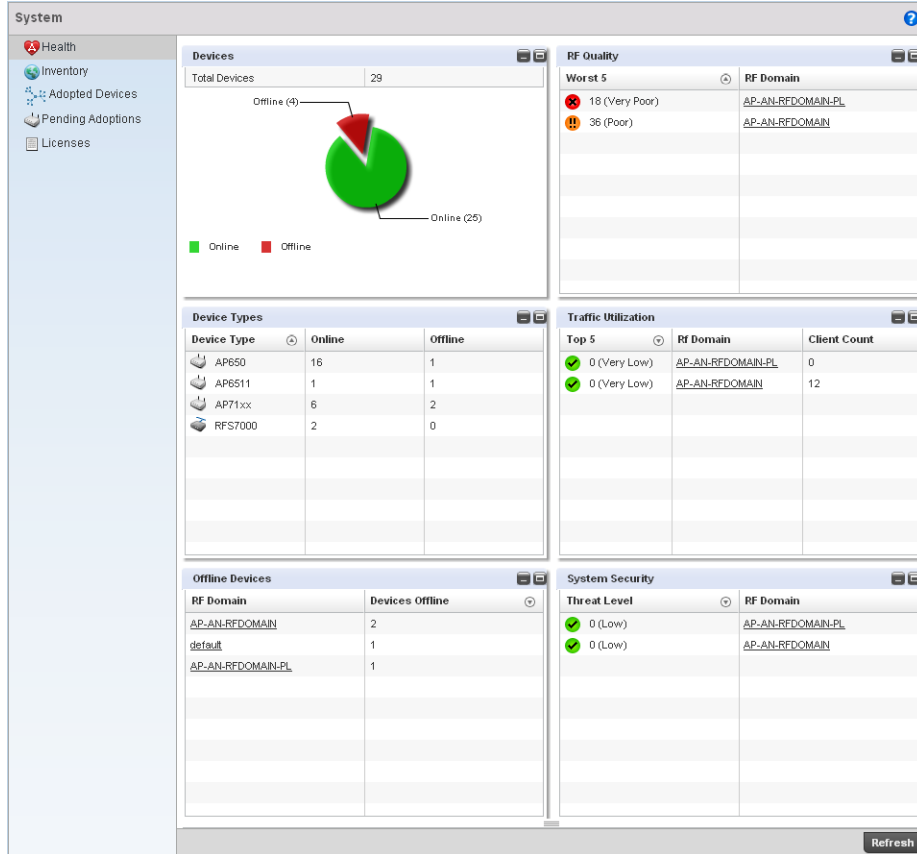


FIGURE 390 System screen

The System screen fields supporting *Device Health*, *Device Type*, *RF Quality Index*, *Traffic Utilization* and *Security*.

4. The **Device Health** table displays the total number of devices in the managed network. The pie chart is a proportional view of how many devices are functional and are currently online. Green indicates online devices and the red offline devices.
5. Use the **RF Quality Index** table to isolate poorly performing RF Domains. This information is a starting point to improving the overall quality of the wireless controller managed network.
6. The **RF Quality Index** table displays the RF Domain RF performance. Quality indices are:
 - 0–50 (Poor)
 - 50–75 (Medium)
 - 75–100 (Good).

This area displays the following:

Worst 5 Displays five RF Domains with the lowest quality indices in the wireless controller-managed network. The value can be interpreted as:

- 0-50 – Poor quality
- 50-75 – Medium quality
- 75-100 – Good quality

RF Domain Displays the name of the RF Domain.

7. The **Traffic Utilization** table displays the top 5 RF Domains with the most effective resource utilization. Utilization is dependent on the number of devices connected to the RF Domain.

Top 5 •

RF Domain Displays the name of the RF Domain.

Client Count Displays the number of wireless clients associated with the RF Domain.

8. The **Security** table defines a Threat Level as an integer value indicating a potential threat to the system. It's an average of the threat indices of all the RF Domains managed by the wireless controller.

Threat Level Displays the threat perception value. This value can be interpreted as:

- 0-2 – Low threat level
- 3-4 – Moderate threat level
- 5 – High threat level

RF Domain Displays the name of the RF Domain for which the threat level is displayed.

Inventory

[System Statistics](#)

The *Inventory* screen displays information about the physical hardware managed by the wireless controller. Use this information to assess the overall performance of managed devices.

To display the inventory statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select the **System** node from the left navigation pane.
3. Select **Inventory** from the left-hand side of the UI.

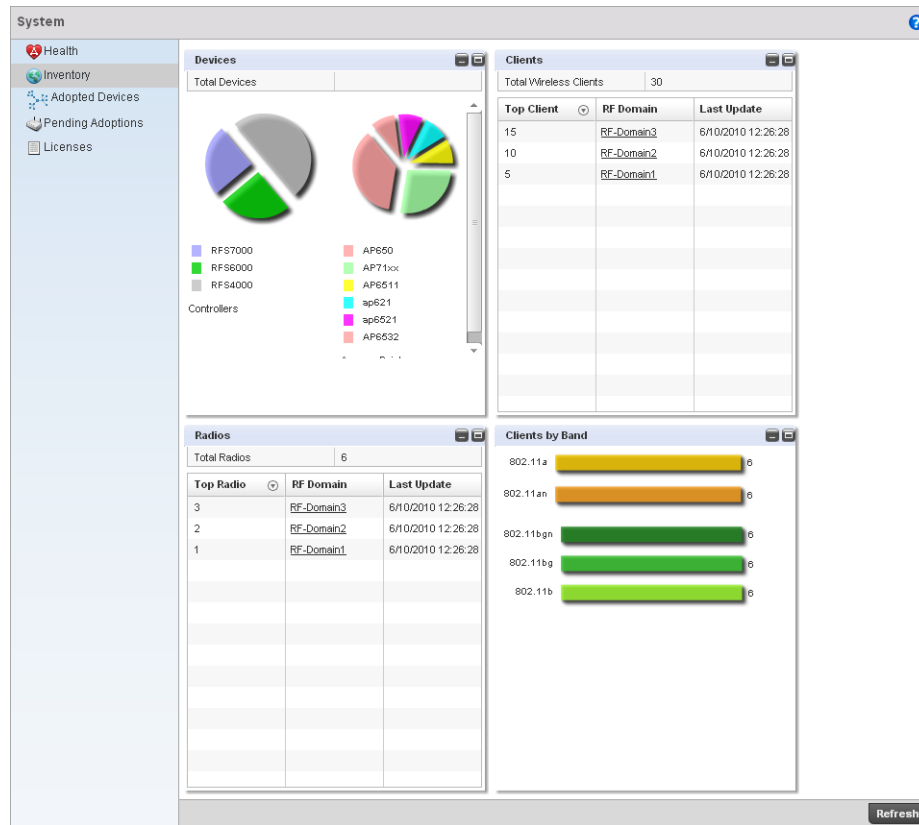


FIGURE 391 System Inventory screen

- The **Device Types** table displays an exploded pie chart depicting device type distribution and members of the managed network. The table in the Device Type area displays the total number of devices managed by this controller.
- The **Radios** table displays radios in use throughout within the wireless controller managed network. This area displays the total number of managed radios and top 5 RF Domains in terms of radio count. The Total Radios value is the total number of radios in this system.

Top Radio Count	Displays the number of radios in the RF Domain.
RF Domain	Displays the name of the RF Domain these radios belong.
Last Update	Displays the UTC timestamp when this value was reported.

- The **Wireless Clients** table displays the total number of wireless clients managed by the wireless controller. This Top Client Count table lists the top 5 RF Domains, in terms of the number of wireless clients adopted:

Top Client Count	Displays the number of wireless clients adopted by the RF Domain.
RF Domain	Displays the name of the RF Domain.
Last Update	Displays the UTC timestamp when the client count was last reported.

- The **Clients on 5 GHz Channels** area displays the number of clients using 5 GHz radios.

8. The **Clients on 2.4 GHz Channels** area displays the number of clients using 2.4 GHz radios.

Adopted Devices

System Statistics

The *Adopted Devices* screen displays a list of devices adopted to the wireless controller managed network. Use this screen to view a list of devices and their current status.

To view adopted AP statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select the **System** node from the left navigation pane.
3. Select **Adopted Devices** from the left-hand side of the UI.

Adopted Device	RF Domain Name	Type	Model Number	Adopted by	Adoption Time	Uptime
00-23-68-0F-44-24	AP-AN-RFDOMAIN	AP71xx	AP7131N-US	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-0F-C5-AC	AP-AN-RFDOMAIN	AP71xx	AP7131	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-0F-C6-F8	AP-AN-RFDOMAIN	AP71xx	AP7131	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-31-1A-08	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-31-1A-94	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-3A-71-F8	AP-AN-RFDOMAIN	AP71xx	AP7131N-US	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-41-DC-50	AP-AN-RFDOMAIN	AP71xx	AP7131N-US	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-85-8F-84	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-85-90-24	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-85-92-80	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-86-44-B4	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-86-44-B8	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-86-45-08	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-86-45-38	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-86-45-44	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-86-45-48	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-86-45-4C	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-86-45-5C	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-86-45-8C	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-86-47-BC	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-86-47-C0	AP-AN-RFDOMAIN	AP650	AP-0650-66030-L	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
00-23-68-9E-51-70	AP-AN-RFDOMAIN	AP71xx	AP7131N-US	00-15-70-37-FB-18	Wed May 11 2011	Wed May 11 2011
5C-0E-8B-08-44-F8	AP-AN-RFDOMAIN-PL	AP6511	AP-6511-60010-L	00-15-70-37-FB-18	Fri May 13 2011 02	Fri May 13 2011 02

FIGURE 392 System Adopted Devices screen

4. The **Adopted Devices** screen provides the following:

Adopted Device	Displays the hostname of the adopted device.
MAC Address	Displays the MAC address of the adopted device.
Type	Displays the AP type (either AP650, AP7131, or AP6511).
RF Domain Name	Displays the domain the adopted AP belongs to.

- Online** Displays the working state of the adopted AP. A red 'X' indicates the device is offline and not working.
- Version** Displays the software version of each listed AP. This information can be helpful when troubleshooting and working with support personnel.
- Serial Number** Displays the AP serial number. The serial number is used for controller management. This field is read-only and can not be modified.

Pending Adoptions

System Statistics

The *Adopted Devices* screen displays a list of devices adopted to the wireless controller managed network. Use this screen to view a list of devices and their current status.

To view adopted AP statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select the **System** node from the left navigation pane.
3. Select **Pending Adoptions** from the left-hand side of the UI.

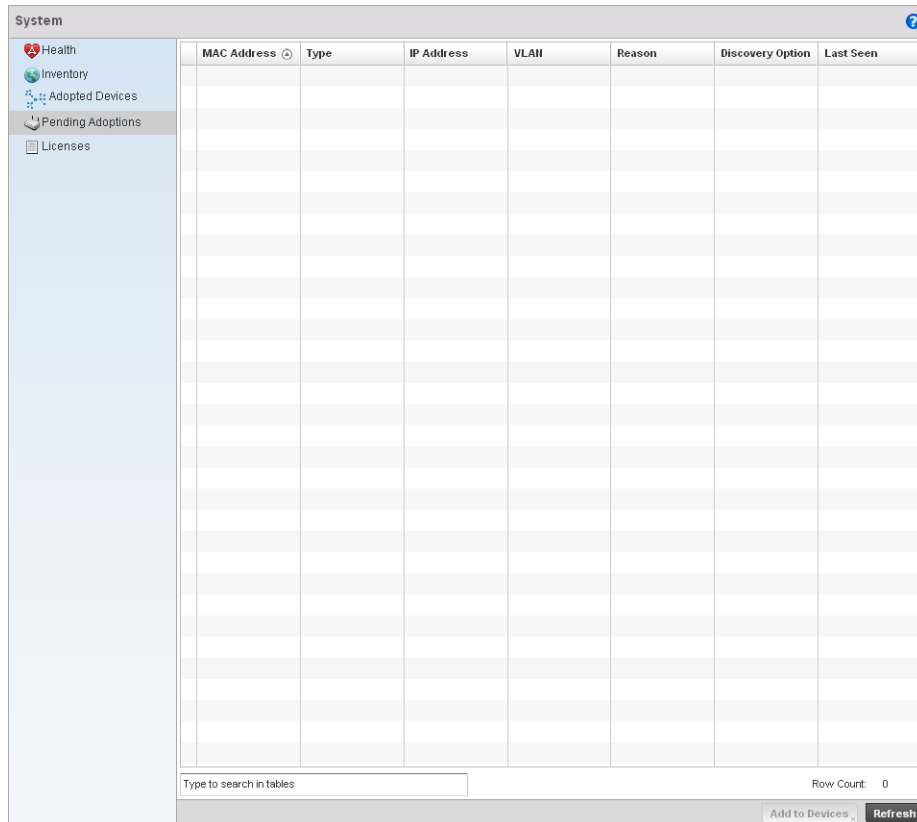


FIGURE 393 Pending Adoptions Devices screen

4. The **Adopted Devices** screen provides the following

MAC Address	Displays the MAC address of the device pending adoption.
Type	Displays the AP type (either AP650, AP7131, or AP6511).
IP Address	Displays the current IP Address of the device pending adoption.
VLAN	Displays the current VLAN number of the device pending adoption.
Reason	Displays the status as to why the device is still pending adoption.
Discovery Option	Displays the discovery option code for each AP listed pending adoption.
Last Seen	Displays the date and time stamp of the last time the device was seen. Click the arrow next to the date and time to toggle between standard time and UTC.

Licenses

[System Statistics](#)

The licenses statistics screen displays available licenses for devices within a cluster. It displays the total number of AP licenses and adaptive AP licenses.

To view a licenses statistics of a managed network:

1. Select the **Statistics** menu from the Web UI.
2. Select the **System** node from the left navigation pane.
3. Select **Licenses** from the left-hand side of the UI.

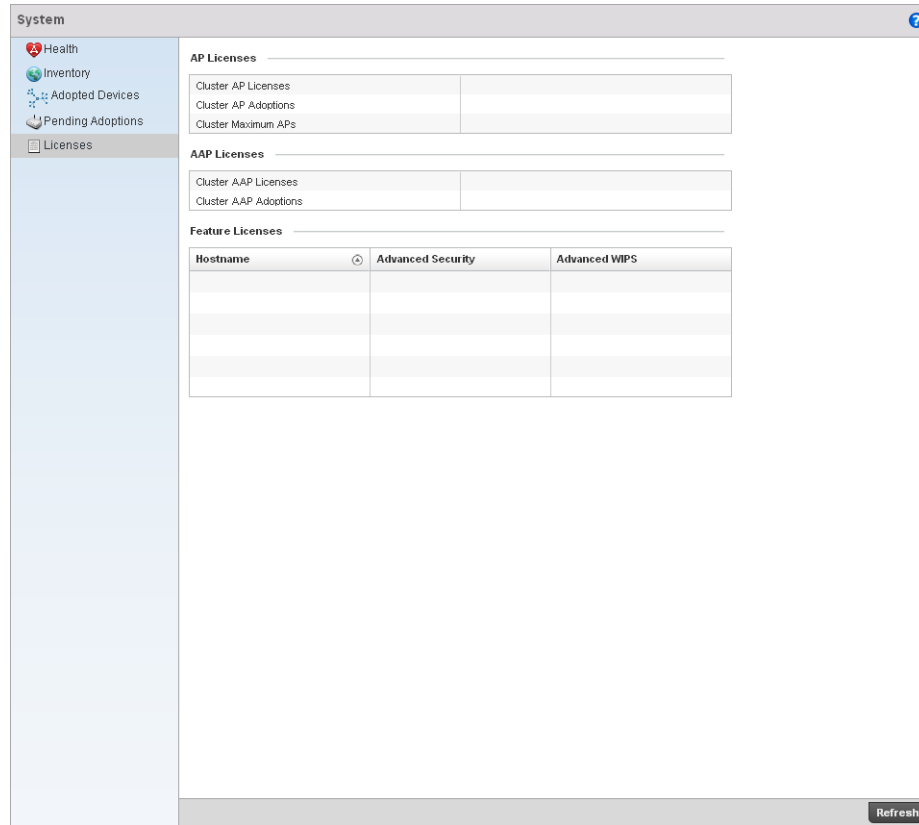


FIGURE 394 System Licenses screen

4. The **AP Licenses** area provides the following information:

- Cluster AP Licenses** Displays the number of access point licenses installed in the cluster.
- Cluster AP Adoptions** Displays the number of Access Points points adopted by the cluster.
- Cluster Maximum APs** Displays the maximum number of Access Points that can be adopted by the controllers in the cluster.

5. The **AAP Licenses** area provides the following information:

- Cluster AAP Licenses** Displays the number of adaptive Access Points in the cluster.
- Cluster AAP Adoptions** Displays the number of Access Points adoptable by the controllers in a cluster.

6. The **Featured Licenses** area provides the following information:

- Hostname** Displays the hostname for each feature license installed.
- Advanced Security** Displays whether the Advanced Security feature is installed for each hostname.
- Advanced WIPS** Displays whether the Advanced WIPS feature is installed for each hostname.

RF Domain Statistics

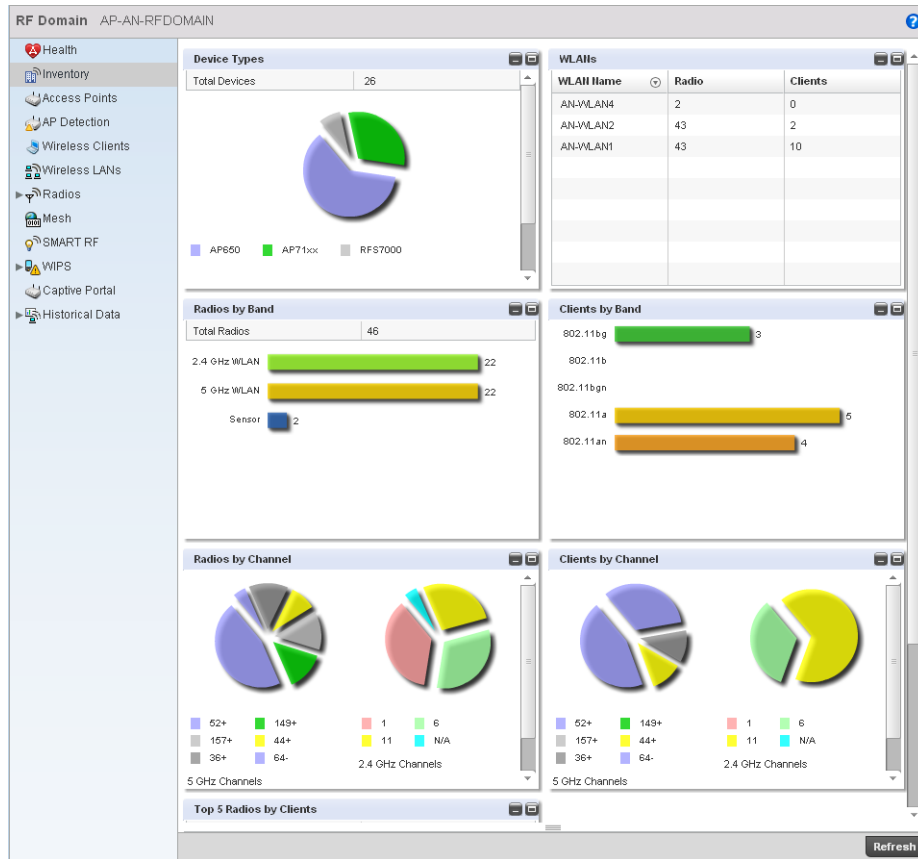
1. Select a **RF Domain** node from the left navigation pane.

The screenshot displays the RF Domain Manager interface for 'RF-Domain1'. The left navigation pane includes: Health, Inventory, Access Points, AP Detection, Wireless Clients, Wireless LANs, Radios, Mesh, SMART RF, WIPS, Captive Portal, and Historical Data. The main content area is divided into several panels:

- Domain:** Shows 'RF Domain Manager'.
- Radio Traffic Utilization:** Displays 'Traffic Index' at 23 (Low) with a warning icon and 'Max: User Rate' at 50,000. Below is a table for 'Top 5' radios.
- Devices:** Shows 'Total Devices' as 3 and a large green circle representing the overall status, with a legend for 'Online'.
- Client Traffic Utilization:** Displays a table for 'Top 5' clients with columns for Client MAC and Vendor.
- Radio Quality:** Shows 'RF Quality Index' at 88 (Good) with a green checkmark. Below is a table for 'Worst 5 Radios'.
- Smart RF Activity:** Lists 'Power Changes', 'Channel Changes', and 'Coverage Changes'.

A 'Refresh' button is located at the bottom right of the interface.

2. statusstatusoverall effectiveness of the RF environment as a percentage of the connect rate in both directions, as well as the retry and error rate. This area also lists the worst 5 radios in the RF Domain. The RF Quality Index can be interpreted as: 0-20 – Very poor quality 20-40 – Poor quality 40-60 – Average quality 60-100 – Good quality Refer to the table below for the fields in the **Worst 5 Radios** table: Refer to the table below for the fields in the **Worst 5 Clients** table: Refer to the table below for the fields in the **Top 5 Radios** table: Refer to the table below for the fields in the **Top 5 Clients** table: Select a **RF Domain** node from the left navigation pane.



3. Select a RF Domain node from the left navigation pane.

RF Domain RF-Domain1

Access Point	AP MAC Address	Type	Client Count	Radio Count
AP1-ControllerA-AP650	AA-11-00-00-00-00	AP650	10	2
AP2-ControllerA-AP71xx	AA-22-00-00-00-00	AP71xx	33	2

Type to search in tables Row Count: 2

[Refresh](#)

- Select a **RF Domain** node from the left navigation pane.

RF Domain RF-Domain1

The screenshot shows a network management interface. On the left is a navigation pane with various icons and labels: Health, Inventory, Access Points, AP Detection, Wireless Clients, Wireless LANs, Radios, Mesh, SMART RF, WIPS, Captive Portal, and Historical Data. The main area displays a table with the following columns: BSSID, Channel, SSID, RSSI, and Reported By. The first row contains the data: Channel 11, SSID evilbit, and RSSI 60 dBm. Below the table is a search input field and a 'Row Count: 1' indicator. At the bottom right are 'Clear All' and 'Refresh' buttons.

BSSID	Channel	SSID	RSSI	Reported By
	11	evilbit	60 dBm	

Type to search in tables

Row Count: 1

Clear All Refresh

5. Select a **RF Domain** node from the left navigation pane.

RF Domain RF-Domain1

MAC Address	WLAN	Username	State	VLAN	IP Address	Vendor
AA-11-11-00-00-00	wlan1	user1	associating	1	10.1.1.1	Motorola
AA-11-22-00-00-00	wlan1	user1	associating	2,100	10.1.1.1	Motorola
AA-11-33-00-00-00	WLAN2				10.1.3.0	
AA-22-11-00-00-00	wlan1	user1	associating	1	10.1.1.1	Motorola
AA-22-11-00-00-00	WLAN2				10.2.1.0	
AA-22-22-00-00-00	wlan1	user1	associating	2,100	10.1.1.1	Motorola
AA-22-22-00-00-00	WLAN2				10.2.2.0	
AA-22-33-00-00-00	WLAN2				10.2.3.0	
AA-33-11-00-00-00	WLAN2				10.3.1.0	
AA-33-22-00-00-00	WLAN2				10.3.2.0	
AA-33-33-00-00-00	WLAN2				10.3.3.0	
AA-33-33-01-00-00	WLAN2				10.3.3.0	
AA-33-33-02-00-00	WLAN2				10.3.3.0	
AA-33-33-03-00-00	WLAN2				10.3.3.0	
AA-33-33-04-00-00	WLAN2				10.3.3.0	
AA-33-33-05-00-00	WLAN2				10.3.3.0	
AA-33-33-06-00-00	WLAN2				10.3.3.0	
AA-33-33-07-00-00	WLAN2				10.3.3.0	
AA-33-33-08-00-00	WLAN2				10.3.3.0	
AA-33-33-09-00-00	WLAN2				10.3.3.0	
AA-33-33-10-00-00	WLAN2				10.3.3.0	

Type to search in tables Row Count: 21

[Refresh](#)

FIGURE 395

6. Select a RF Domain node from the left navigation pane.

RF Domain RF-Domain1

WLAN Name	SSID	Traffic Index	Radio Count	Tx Bytes	Tx User Data Rate	Rx Bytes	Rx User Data Rate
WLAN1	WLAN1	71 (High)	2	300	400 kbps	200	100 kbps
WLAN2	WLAN2	99 (High)	3	2,300	4,100 kbps	2,100	1,001 kbps

Type to search in tables Row Count: 2

Refresh

7. Select a RF Domain node from the left navigation pane.

RF Domain RF-Domain1

Radio	Radio MAC	Radio Type	State	Channel Current (Config)	Power Current (Config)	Clients
ap6511-9a:6c:R1	11:22:33:44:55:66	2.4 GHz WLAN	s5	5 (7)	15 (15)	
ap6511-9b:6c:R2	22:22:33:44:55:66	5 GHz WLAN	s5	60 (100)	17 (16)	

Type to search in tables Row Count: 2

Refresh

8.

RF Domain RF-Domain1

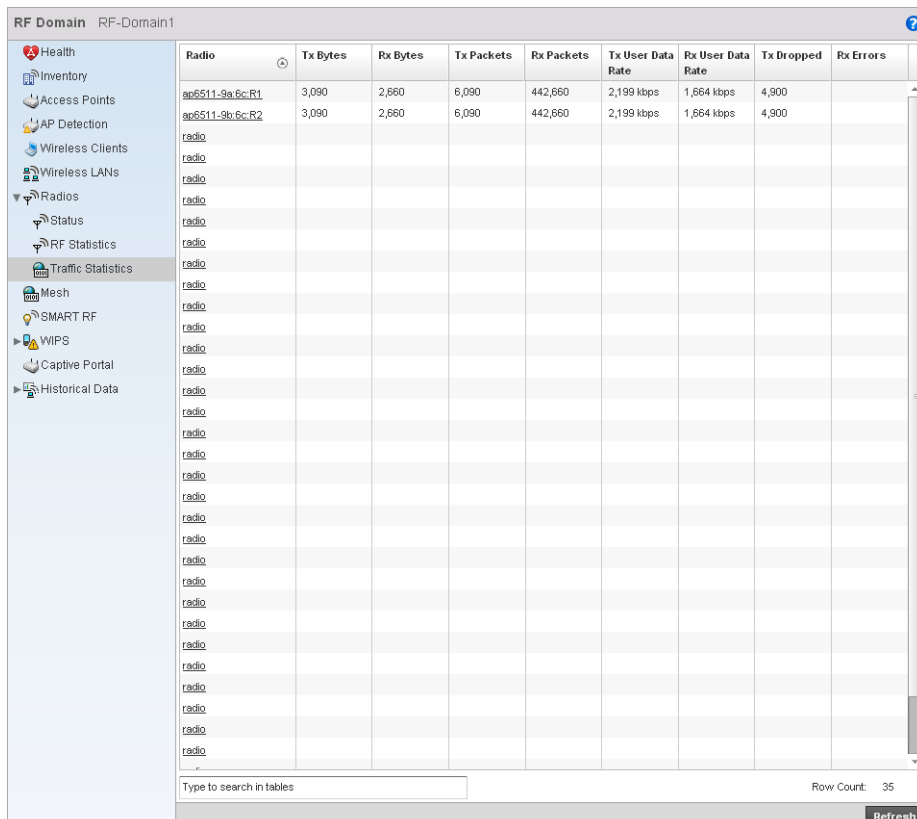
- Health
- Inventory
- Access Points
- AP Detection
- Wireless Clients
- Wireless LANs
- Radios
- Status
- RF Statistics
- Traffic Statistics
- Mesh
- SMART RF
- WIPS
- Captive Portal
- Historical Data

Radio	Signal	SIIR	Tx Physical Layer Rate	Rx Physical Layer Rate	Avg. Retry Number	Error Rate	Traffic Index	RF Quality Index
ap6511-9a.6c:R1	455 dbm	344 db	34,955 Mbps	4,459 Mbps	24	3	34 (Poor)	
ap6511-9b.6c:R2	455 dbm	344 db	34,955 Mbps	4,459 Mbps	24	38	35 (Poor)	
radio							18 (Very Poor)	
radio							84 (Good)	
radio							92 (Good)	
radio							100 (Good)	
radio							94 (Good)	
radio							95 (Good)	
radio							92 (Good)	
radio							100 (Good)	
radio							95 (Good)	
radio							87 (Good)	
radio							100 (Good)	
radio							100 (Good)	
radio							97 (Good)	
radio							97 (Good)	
radio							99 (Good)	
radio							92 (Good)	
radio							97 (Good)	
radio							91 (Good)	
radio							70 (Good)	
radio							91 (Good)	
radio							99 (Good)	
radio							100 (Good)	
radio							100 (Good)	
radio							97 (Good)	
radio							100 (Good)	
radio							91 (Good)	
radio							85 (Good)	
radio							100 (Good)	
radio							93 (Good)	
radio							90 (Good)	
radio							90 (Good)	

Type to search in tables Row Count: 35

Refresh

9. Select a **RF Domain** node from the left navigation pane.



The screenshot displays the 'RF Domain' interface. On the left is a navigation pane with categories including Health, Inventory, Access Points, AP Detection, Wireless Clients, Wireless LANs, Radios (expanded to show Status and RF Statistics), Traffic Statistics, Mesh, SMART RF, WIPS, Captive Portal, and Historical Data. The 'RF Statistics' table is active, showing two rows of data for radio nodes. The table has columns for Radio, Tx Bytes, Rx Bytes, Tx Packets, Rx Packets, Tx User Data Rate, Rx User Data Rate, Tx Dropped, and Rx Errors. The first two rows show identical data: 3,090 Tx/Rx Bytes, 6,090 Tx/Rx Packets, 2,199 kbps Tx/Rx User Data Rates, and 4,900 Tx Dropped. The remaining rows are labeled 'radio' and are currently empty. At the bottom, there is a search bar and a 'Refresh' button. The row count is displayed as 'Row Count: 35'.

Radio	Tx Bytes	Rx Bytes	Tx Packets	Rx Packets	Tx User Data Rate	Rx User Data Rate	Tx Dropped	Rx Errors
ap6511-9a.6c.R1	3,090	2,660	6,090	442,660	2,199 kbps	1,664 kbps	4,900	
ap6511-9b.6c.R2	3,090	2,660	6,090	442,660	2,199 kbps	1,664 kbps	4,900	
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								
radio								

The screenshot shows a web-based network management interface. On the left is a navigation pane with a tree view containing the following items: Health, Inventory, Access Points, AP Detection, Wireless Clients, Wireless LANs, Radios (expanded), Status, RF Statistics, Traffic Statistics, Mesh (selected), SMART RF, WIPS, Captive Portal, and Historical Data. The main area displays a table with the following columns: Client, Client Radio MAC, Portal, Portal Radio MAC, and Connect Time. The table is currently empty. At the bottom of the table area, there is a search input field with the placeholder text 'Type to search in tables', a 'Row Count: 0' indicator, and a 'Refresh' button.

Client	Client Radio MAC	Portal	Portal Radio MAC	Connect Time
--------	------------------	--------	------------------	--------------

10. Select a **RF Domain** node from the left navigation pane.

The screenshot displays the Brocade Mobility management interface for an RF Domain. The left navigation pane includes sections like Health, Inventory, Access Points, AP Detection, Wireless Clients, Wireless LANs, Radios, Status, RF Statistics, Traffic Statistics, Mesh, SMART RF, WIPS, Captive Portal, and Historical Data. The SMART RF section is expanded, showing 'SMART RF Details' with two MAC addresses: 11:22:33:44:55:66 and 41:22:33:44:55:66. The main area shows a table titled 'Radios' with columns for AP MAC Address, MAC Address, Type, State, Channel, and Power. Two rows of data are visible in the table.

AP MAC Address	MAC Address	Type	State	Channel	Power
11:22:33:44:55:66	11:22:33:44:55:66	344	s2	9	11 dBm
41:22:33:44:55:66	41:22:33:44:55:66	3441	s21	91	21 dBm

WIPS

11. Select a **RF Domain** node from the left navigation pane.

The screenshot displays the 'RF Domain' interface for 'RF-Domain1'. On the left is a navigation pane with various RF-related tools. The main area shows a table titled 'Client Blacklist' with two rows of data.

Event Name	Blacklisted Client	Time Blacklisted	Total Time	Time Left
dos-eapol-start-storm	44-55-44-55-44-55	Thu Jun 10 2010 12:26:28 PM	2h 0m 0s	1h 0m 0s
null-probe-response	44-55-44-55-44-55	Thu Jun 10 2010 12:26:28 PM	40m 0s	20m 0s

At the bottom of the interface, there is a search box labeled 'Type to search in tables' and a 'Row Count: 2' indicator. A 'Refresh' button is located in the bottom right corner.

12. Select a **RF Domain** node from the left navigation pane.

The screenshot shows the 'RF Domain' interface for 'RF-Domain1'. The left navigation pane includes categories like Health, Inventory, Access Points, AP Detection, Wireless Clients, Wireless LANs, Radios, Status, RF Statistics, Traffic Statistics, Mesh, SMART RF, WIPS, Client Blacklist, WIPS Events (highlighted), Captive Portal, and Historical Data. The main content area is a table with the following data:

Event Name	Reporting AP	Originating Device	Detector Radio	Time Reported
dos-eapol-start-storm	AP1-ControllerA-AP650	33-44-33-44-33-44	1	Thu Jun 10 2010 12:26:28 PM
null-probe-response	AP1-ControllerA-AP650	33-44-33-44-33-44	1	Thu Jun 10 2010 12:26:28 PM

At the bottom of the table, there is a search input field with the placeholder 'Type to search in tables', a 'Row Count: 2' indicator, and two buttons: 'Clear All' and 'Refresh'.

13. Select a **RF Domain** node from the left navigation pane.

RF Domain RF-Domain1 ?

Client MAC	Client IP	Captive Portal	Authentication	VLAN	VLAN	Remaining Time
AA-11-11-00-00-00	1.1.1.1	default	Success	VLAN3	1	1m 40s
AA-11-12-00-00-00	1.1.1.1	default	Pending	VLAN4	2	3m 20s

Type to search in tables Row Count: 2

Refresh

Viewing Smart RF History

Historical Data

To view Smart RF history:

1. Select the **Statistics** menu from the Web UI.
2. Select a **RF Domain** node from the left navigation pane.
3. Select **Historical Data** > **SMART RF History** from the left-hand side of the UI.

AP MAC	Radio MAC	Radio Index	Type	New Value	Old Value	Time
aa:aa:aa:aa:aa:aa	11:22:33:44:55:66	R2	AP Unadopted	5	6	
bb:aa:aa:aa:aa:aa	12:22:33:44:55:66	R3	AP Unadopted	55	65	
c:aa:aa:aa:aa:aa	13:22:33:44:55:66	R4	AP Unadopted	45	46	

FIGURE 396 RF Domain Smart RF History

This screen displays the following:

AP MAC	Displays the MAC address of the selected AP.
Radio MAC	Displays the radio MAC address of the corresponding AP.
Radio Index	Displays the numerical identifier assigned to each detector AP used in calibration.
Type	Displays the AP type.
New Value	Displays the new power value as assigned by Smart RF.
Old Value	Lists the old power value before calibration resulted in modifications.
Time	Displays time stamp when this event occurred.

Access Point Statistics

The Access Point Statistics screen displays APs available within the managed network. Use this data as necessary to check all whether APs are active, their VLAN assignments and their current authentication and encryption schemes. Access Point Statistics consists of the following:

- [Health](#)
- [Inventory](#)

- [Device](#)
- [AP Upgrade](#)
- [Wireless Client](#)
- [Wireless LANs](#)
- [Radios](#)
- [Mesh](#)
- [Mesh](#)
- [Network](#)
- [Firewall](#)
- [Certificates](#)
- [WIPS](#)
- [Sensor Servers](#)
- [Captive Portal](#)

Health

[Access Point Statistics](#)

The *Health* screen displays information on the selected device, such as its hardware and software version. Use this information to fine tune the performance of the selected APs. This screen should also be the starting point for troubleshooting.

To view the access point health:

Select the **Statistics** menu from the Web UI.

Select an **Access Point** node from the left navigation pane.

Select **Health** from the left-hand side of the UI.

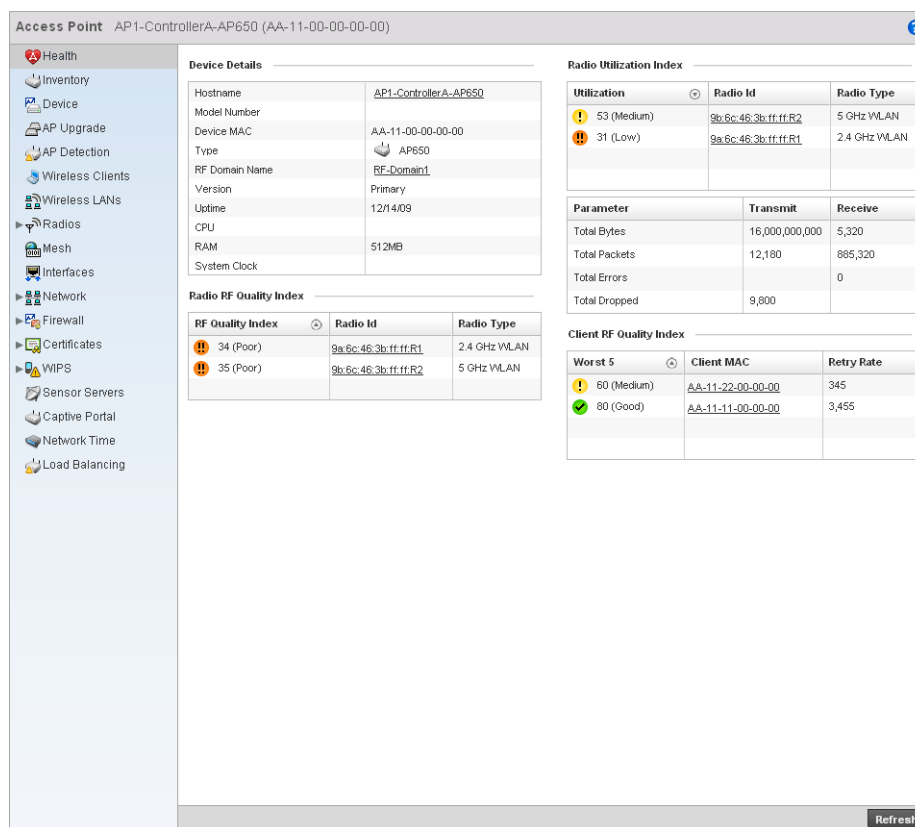


FIGURE 397 Access Point Health screen

The **Device Details** area displays the following:

Hostname	Displays the AP's unique name. A hostname is assigned to a device connected to a computer network.
Device MAC	Displays the MAC address of the AP. This is factory assigned and cannot be changed.
Type	Displays the Access Point's model.
RF Domain Name	Displays an AP's RF Domain membership.
Version	Displays the AP's current firmware version. Use this information to assess whether an upgrade is required for better compatibility with the controller.
Uptime	Displays the cumulative time since the AP was last rebooted or lost power.
CPU	Displays the processor core.
RAM	Displays the free AP memory available.
System Clock	Displays system clock information.

The **RF Quality Index** table displays the following:

RF Quality Index	Displays radios with very low quality indices. RF quality index indicate overall RF performance. The RF quality indices are: <ul style="list-style-type: none"> • 0-50 (poor) • 50-75 (medium) • 75-100 (good)
Radio ID	Displays a radio's hardware encoded MAC address.
Radio Type	Identifies whether the radio is a 802.11b, 802.11bg, 802.11bgn, 802.11a, or 802.11an.

The **Radio Utilization Index** tables display the following:

Utilization	Displays the traffic indices of radios, which measures how efficiently the traffic medium is used. This value is indicated as an integer.
Radio Id	Displays a numerical value assigned to the radio as a unique identifier. For example: 1, 2, or 3.
Radio Type	Identifies whether the radio is an 802.11b, 802.11bg, 802.11bgn, 802.11a, or an 802.11an.
Parameter	Displays the statistics in number of packets for: <ul style="list-style-type: none"> • Total Bytes - The total number of bytes that passed through the Access Point. • Total Packets - The total number of packets that passed through the Access Point. • Total Errors - The total error packets. • Total Dropped - The total dropped packets.
Transmit	Displays the total number of packets transmitted by the radio.
Receive	Displays the total number of packets received by the radio.

The **Client RF Quality Index** table displays the following:

Worst 5	Displays the 5 clients having the lowest RF quality.
Client MAC	Displays the MAC address of the client with low RF indices.
Retry Rate	Displays the average number of retries per packet. A high number indicates potential network or hardware problems.

Inventory

[Access Point Statistics](#)

The *Inventory* screen displays information about the hardware managed by the wireless controller. Use the information provided in this screen to understand the overall performance of managed devices.

The *Inventory* screen also displays information about AP physical characteristics. Use this screen to gather information on the performance of the different clients associated with the AP. Additionally, use this screen to fine tune Access Point performance.

To view the access point inventory statistics:

Select the **Statistics** menu from the Web UI.

Select an **Access Point** node from the left navigation pane.

Select **Inventory** from the left-hand side of the UI.

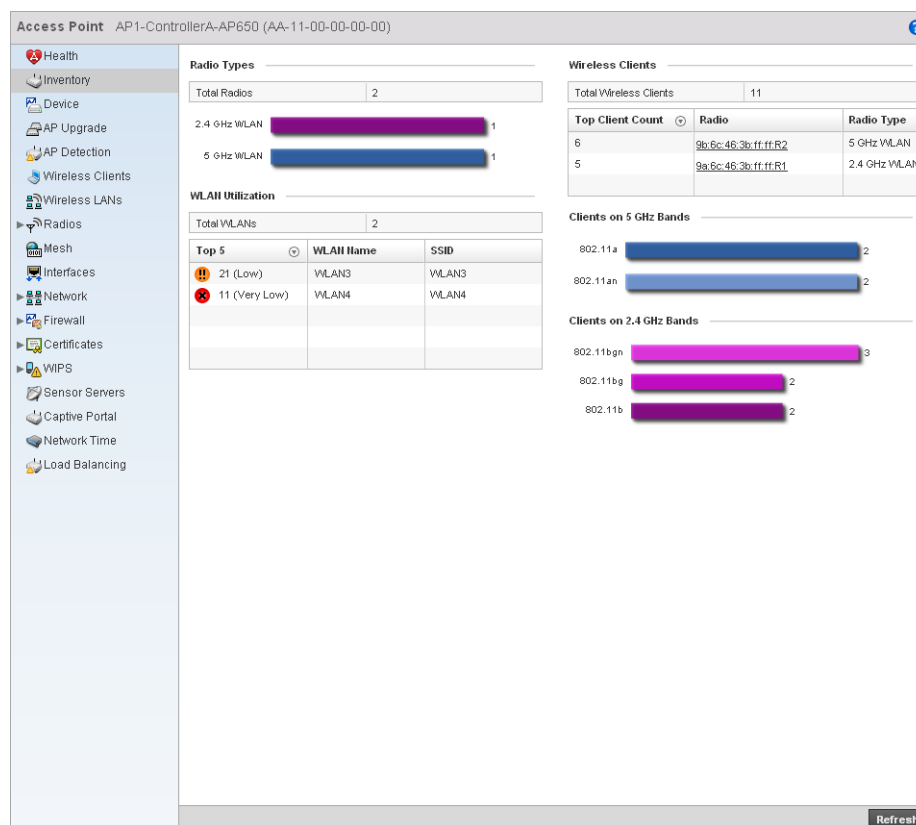


FIGURE 398 Access Point Inventory screen

The **Radio Types** table displays the total number of radios detected. It also displays the number of radios using the 2.4 GHz and 5 GHz frequency bands.

The **Wireless LANs** area displays the total number of WLANs. It also displays the following:

Top 5	Displays the maximum traffic utilization of the WLAN in which the access point is a member. The integer denotes the traffic index, which measures how efficiently the traffic medium is used. Traffic indices are: <ul style="list-style-type: none"> • 0 – 20 (very low) • 20 – 40 (low) • 40 – 60 (moderate) • 60 and above (high)
WLAN Name	Displays a name assigned to identify the WLAN.
SSID	Displays the Service Set ID associated with the WLAN.

The **Wireless Clients** area displays the total number of wireless clients associated with this Access Point. It also displays the following:

Top Client Count	Displays the number of clients associated with this Access Point.
Radio	Displays the radio MAC address associated with the client.
Radio Type	Displays the radio type.

The **Clients on 5 GHz Channels** table displays the number of wireless clients with radios operating in the 5 GHz frequency band.

The **Clients on 2.4 GHz Channels** area displays the number of wireless clients with radios operating in the 2.4 GHz band.

Device

[Access Point Statistics](#)

The **Device** screen displays basic information about the selected Access Point. Use this screen to assess the version, boot image and upgrade status.

To view the device statistics:

Select the **Statistics** menu from the Web UI.

Select an **Access Point** node from the left navigation pane.

Select **Device** from the left-hand side of the UI.

The screenshot shows the 'Access Point Device' screen for an AP1-ControllerA-AP650. The left navigation pane includes options like Health, Inventory, Device, AP Upgrade, AP Detection, Wireless Clients, Wireless LANs, Radios, Mesh, Interfaces, Network, Firewall, Certificates, WIPS, Sensor Servers, Captive Portal, Network Time, and Load Balancing. The main content area is divided into several sections:

- System:** A table with the following data:

Model Number	
Version	Primary
Boot Partition	Primary
Fallback Enabled	✓ Enabled
Fallback Image Triggered	✓ True
Next Boot	Secondary
- Firmware Images:** A table with the following data:

Primary Build Date	12/14/09
Primary Install Date	12/14/09
Primary Version	Primary
Secondary Build Date	12/14/09
Secondary Install Date	12/14/09
Secondary Version	Secondary
- Upgrade Status:** A table with the following data:

Upgrade Status	Successful
Upgrade Status Time	Jan 13
- DNS Mappings:** A table with columns 'Name Server' and 'Type'. It is currently empty.

A 'Refresh' button is located at the bottom right of the screen.

FIGURE 399 Access Point Device screen

The **System** table displays the following:

Model Number	Displays the Access Point model number for the selected AP.
Version	Displays the software (firmware) version on the access point.
Boot Partition	Displays the boot partition type.
Fallback Enabled	Displays whether this option is enabled. This method enables a user to store both a known legacy firmware version and a new firmware version in device memory. The user can test the new software, and use an automatic fallback, which loads the old version if the new version fails.
Fallback Image Triggered	Displays whether the fallback image was triggered. The fallback image is an old version of a known and operational software stored in device memory. This allows a user to test a new software version of software. If the new version fails, the user can fall back to the old version.
Next Boot	Designates this version as the version used the next time the AP is booted.

The **Firmware Images** table displays the following:

Primary Build Date	Displays the build date when the version was created.
Primary Install Date	Displays the date this version was installed.
Primary Version	Displays the primary version string.
Secondary Build Date	Displays the build date when this version was created.
Secondary Install Date	Displays the date this secondary version was installed.
Secondary Version	Displays the secondary version string.

The **Upgrade Status** table displays the following:

Upgrade Status	Displays the status of the image upgrade.
Upgrade Status Time	Displays the time of the image upgrade.

The **DNS Mappings** table displays the following:

Name Server	Displays any custom Name Server mappings on the Access Point.
Type	Displays the type of DNS mapping, if any, on the Access Point.

AP Upgrade

[Access Point Statistics](#)

To view the AP Upgrade statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **AP Upgrade** from the left-hand side of the UI.

2. Select an **Access Point** node from the left navigation pane.
3. Select **AP Detection** from the left-hand side of the UI.

Unsanctioned AP	Reporting AP	SSID	AP Mode	Radio Type	Channel	RSSI	Last Seen
11:22:33:44:55	AP1-ControllerA-AP650	evlbit	Ad Hoc	11a	11	60 dBm	10s

FIGURE 401 Access Point AP Detection Screen

4. The **AP Detection** screen displays the following:

Unsanctioned	Displays the MAC address of the unsanctioned AP.
Reporting AP	Displays the numerical value for the radio used with the detecting AP.
SSID	Displays the SSID of the WLAN to which the unsanctioned AP belongs.
AP Mode	Displays the mode of the unsanctioned AP.
Radio Type	Displays the type of the radio on the unsanctioned AP. The radio can be 802.11b, 802.11bg, 802.11bgn, 802.11a or 802.11an.
Channel	Displays the channel the unsanctioned AP is currently transmitting on.
Last Seen	Displays the time (in seconds) the unsanctioned AP was last seen on the network by the detecting AP.

Wireless Client

[Access Point Statistics](#)

The *Wireless Clients* screen displays read only device information for wireless clients associated with the selected Access Point. Use this information to assess if configuration changes are required to improve network performance.

To view wireless client statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Wireless Clients** from the left-hand side of the UI.

Client MAC	WLAN	Username	State	VLAN	IP Address	Vendor
AA-11-11-00-00-00	wlan1	user1	associating	1	10.1.1.1	Motorola
AA-11-22-00-00-00	wlan1	user1	associating	2,100	10.1.1.1	Motorola

Type to search in tables Row Count: 2

Refresh

FIGURE 402 Access Point Wireless Clients screen

4. The **Wireless Clients** screen displays the following:

Client MAC	Displays the MAC address of the wireless client.
WLAN	Displays the name of the WLAN the client is currently associated with. Use this information to determine if the client/WLAN placement best suits intended operation and the client coverage area.
Username	Displays the unique name of the administrator or operator.
State	Displays the working state of the client.

- VLAN** Displays the VLAN ID the client is currently mapped to.
- IP Address** Displays the unique IP address of the client. Use this address as necessary throughout the applet for filtering, device intrusion recognition, and approval.
- Vendor** Displays the name of the vendor.

Wireless LANs

Access Point Statistics

The *Wireless LAN* statistics screen displays an overview of Access Point WLANs. This screen displays the WLAN names, their SSIDs, traffic utilization, number of radios etc.

To view the wireless LAN statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Wireless LANs** from the left-hand side of the UI.

Access Point AN-01-0FC5AC (00-23-68-0F-C5-AC)

WLAN Name	SSID	Traffic Index	Radio Count	Tx Bytes	Tx User Data Rate	Rx Bytes	Rx User Data Rate
AN-WLAN1	stcwlfb	0 (Very Low)	2	686,162,531	0 kbps	73,335,207	0 kbps
AN-WLAN2	motorola-guest	0 (Very Low)	2	1,783	0 kbps	10,485	0 kbps
AN-WLAN4	7131meshbackha	0 (Very Low)	1	0	0 kbps	0	0 kbps

Type to search in tables Row Count: 3

Refresh

FIGURE 403 Access Point Wireless LANs screen

4. The **Wireless LANs** screen displays the following:

WLAN Name	Displays the name of the WLAN the Access Point is currently associated with.
SSID	Displays the Service Set ID of the WLAN to which the access point is associated.
Traffic Index	<p>Displays the traffic utilization index, which measures how efficiently the traffic medium is used. It's defined as the percentage of current throughput relative to maximum possible throughput. Traffic indices are:</p> <ul style="list-style-type: none"> • 0–20 (very low utilization) • 20–40 (low utilization) • 40–60 (moderate utilization) • 60 and above (high utilization)
Radio Count	Displays the number of radios associated with this WLAN.
Tx Bytes	Displays the average number of transmitted bytes sent on the selected WLAN.
Tx User Data Rate	Displays the transmitted user data rate in kbps.
Rx Bytes	Displays the average number of packets (in bytes) received on the selected WLAN.
Rx User Data Rate	Displays the user's data rate for received packets.

Radios

[Access Point Statistics](#)

The *Radio* screen displays information on Access Point radios. The actual number of radios depend on the Access Point model and type. This screen displays information on a per radio basis. Use this information to refine and optimize the performance of each radio and therefore improve controller network performance.

The Access Point radio statistics screen provides details about associated radios. It provides radio ID, radio type, RF quality index etc. Use this information to assess the overall health of radio transmissions and access point placement.

To view the radio statistics of an access point:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Radios** from the left-hand side of the UI.

Access Point AP1-ControllerA-AP650 (AA-11-00-00-00)

Radio	Radio MAC	Radio Type	State	Channel Current(Config)	Power Current(Config)	Clients
ap6511-9a:6c:R1	11:22:33-44:55:66	2.4 GHz WLAN	s5	1 (1)	7 (11)	
ap6511-9b:6c:R2	22:22:33-44:55:66	5 GHz WLAN	s54	5 (5)	12 (12)	

Type to search in tables Row Count: 2 Refresh

FIGURE 404 Access Point Radios screen

4. This screen provides the following information:

Radio	Displays the model and numerical value assigned to the radio as its unique identifier.
Radio MAC	Displays the MAC address assigned to the radio as its unique hardware identifier.
Radio Type	Defines whether the radio is a 802.11b, 802.11bg, 802.11bgn, 802.11a or 802.11an.
State	Displays the current operational state of each radio.
Channel Current (Config)	Displays the current channel for each radio and the configured channel in parentheses.
Power Current (Config)	Displays the current power level for each radio and the configured power level in parentheses.
Clients	Displays the number of wireless clients associated with the radio.

Access Point AP1-ControllerA-AP650 (AA-11-00-00-00-00) ?

- Health
- Inventory
- Device
- AP Upgrade
- AP Detection
- Wireless Clients
- Wireless LANs
- Radios
 - Status
 - RF Statistics
- Traffic Statistics
- Mesh
- Interfaces
- Network
- Firewall
- Certificates
- WIPS
- Sensor Servers
- Captive Portal
- Network Time
- Load Balancing

Radio	Signal	SIIR	Tx Physical Layer Rate	Rx Physical Layer Rate	Avg. Retry Number	Error Rate	Traffic Index	Quality Index
ap6511-9a.6c:R1	455 dbm	10 db	34,955 Mbps	4,459 Mbps		24	31	!! 34 (Poor)
ap6511-9b.6c:R2	455 dbm	5 db	34,955 Mbps	4,459 Mbps		24	53	!! 35 (Poor)

Row Count: 2

Refresh

Access Point AP1-ControllerA-AP650 (AA-11-00-00-00)

Radio	Tx Bytes	Rx Bytes	Tx Packets	Rx Packets	Tx User Data Rate	Rx User Data Rate	Tx Dropped	Rx Errors
ap6511-9a-6c-R1	8,000,000,000	2,660	6,090	442,660	2,199 kbps	1,664 kbps	4,900	
ap6511-9b-6c-R2	8,000,000,000	2,660	6,090	442,660	2,199 kbps	1,664 kbps	4,900	

Type to search in tables Row Count: 2

[Refresh](#)

5.

Detailed Radio Statistics

Radios

Click on the *Radio Id* to view the detailed radio statistics. This screen provides information such as radio MAC address, configuration, quality of RF, traffic index, radio WLAN, etc.

This screen is partitioned into:

- [Details](#)
- [Traffic](#)
- [WMM TSPEC](#)
- [Wireless LANs](#)
- [AP Radio Graph](#)

Details

Detailed Radio Statistics

To view the detailed radio statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.

3. Select **Radios** from the left-hand side of the UI, and select **Radio Id**.
4. Select the **Details** sub-menu from the left navigation pane of the resulting screen.

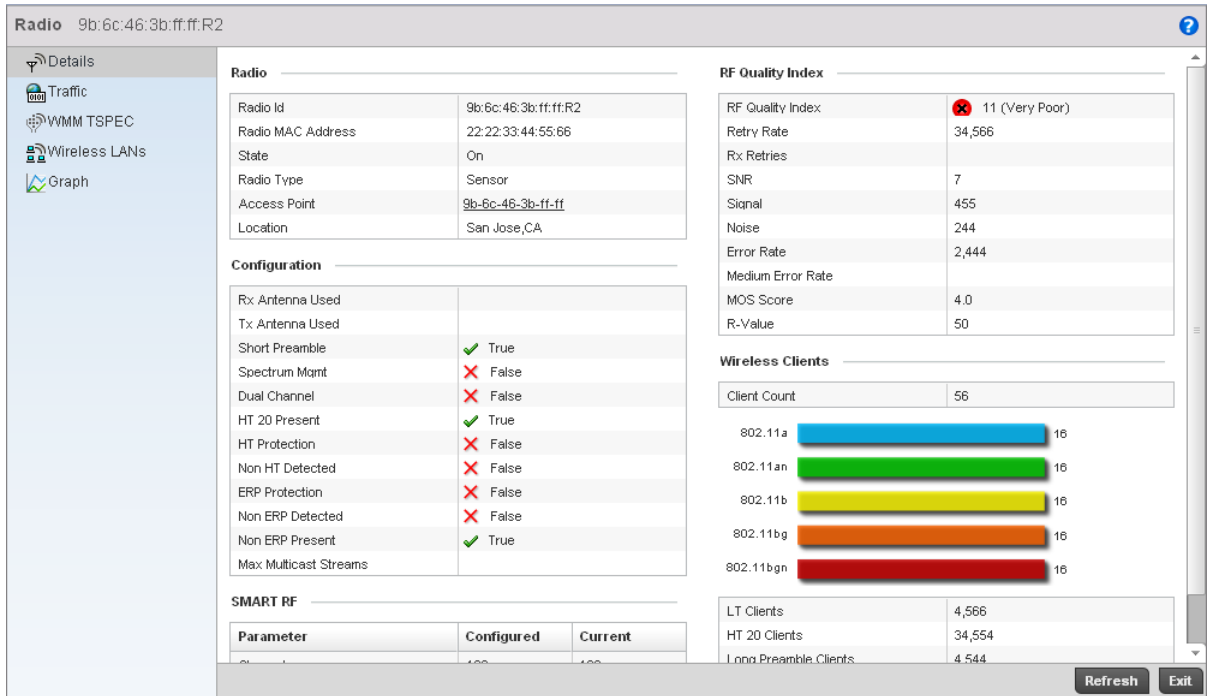


FIGURE 405 Access Point Radio Statistics Details screen

5. The **Radio** table displays the following:

- Radio Id** Displays the numerical index (device identifier) used with this radio.
- Radio MAC Address** Displays the hardware address of the radio. This is read-only and cannot be modified.
- State** Displays the current operational mode of the radio. They can be calibrate, normal, sensor or offline.
- Radio Type** Displays whether the radio is an 802.11b, 802.11bg, 802.11bgn, 802.11a or 802.11an.
- Access Point** Displays the AP this radio belongs to.
- Location** Displays the location of the radio.

6. The **Configuration** table displays the following:

- Rx Antenna Used** Displays the number of receiving antennas supporting this radio.
- Tx Antennas Used** Displays the number of transmitting antennas supporting this radio.
- Short Preamble** Lists whether this feature is enabled on this radio. The radio preamble is a section of data in the packet header that contains information the access point and client need when sending and receiving. A short preamble improves throughput performance. An AP and a client negotiate the use of the short preamble.

Spectrum Mgmt	Displays whether spectrum management is enabled. The purpose of spectrum management is to mitigate radio spectrum pollution and maximize the benefit of the usable radio spectrum.
Dual Channel	Displays whether the radio is transmitting on two channels.
HT 20 Present	Displays whether <i>High Throughput</i> (HT) 20 protection is adopted. In HT 20, all stations in the BSS must be HT stations and must be associated to a 20/40 MHz AP. 20/40 capable HT stations must use protection when transmitting on a 40 MHz channel to prevent the 20 MHz-only HT stations from transmitting at the same time.
HT Protection	Displays whether this feature is enabled. To ensure backward compatibility with older 802.11abg radios, HT access points signal to HT radios to define when to use the protection mode. A field in the beacon frame, called the HT protection field, has 4 possible settings between 0 and 3.
Non HT Detected	Displays whether the AP and clients use HT protection.
ERP Protection	Displays whether this feature is enabled or not. <i>Ethernet Ring Protection</i> (ERP) is a network protection mechanism that enables carriers to enjoy a SONET-like ring protection while leveraging the cost effectiveness of the Ethernet technology. This is an effort to provide sub-50ms protection and recovery switching for Ethernet traffic in a ring topology. This also ensures that there are no loops formed at the Ethernet layer.
Non ERP Detected	Displays whether overlapping networks (not capable of using 802.11g) are detected.
Non ERP Present	Displays when non-802.11g station associates to a network. 802.11g defined the <i>extended rate PHY</i> (ERP). To provide backward compatibility, the ERP information element was defined and is a three-bit flag in a single byte.

7. The **SMART RF** table displays the following:

Channel	The channel list on which Smart RF is enabled.
Power	Lists the power in use for supporting Smart RF on selected channels.

8. The **RF Quality Index** table displays the following:

RF Quality Index	Displays information on the RF quality for the selected wireless client. The RF quality index is the overall effectiveness of the RF environment as a percentage of the connect rate in both directions, as well as the retry and error rate. The RF quality index value can be interpreted as: <ul style="list-style-type: none"> • 0–20 – very poor quality • 20–40 –poor quality • 40–60 –average quality • 60–100 – good quality
Retry Rate	Displays the average number of retries per packet. A high number indicates possible network or hardware problems.
SNR	Displays the <i>signal to noise</i> (SNR) ratio for all the clients connected to selected radio. The SNR is an indication of overall RF performance.
Signal	Displays the power of radio signals in dBm.
Noise	Displays the disturbing influences on the signal by the interference of signals in dBm.

Error Rate	Displays the number of received bit rates altered due to noise, interference and distortion. It's a unitless performance measure.
MOS Score	Displays the average call quality using the <i>Mean Opinion Score</i> (MOS) call quality scale. The MOS scale rates define quality in a scale of 1-5, with higher scores being better. If the MOS score is lower than 3.5, it's likely users will not be satisfied with the voice quality of their calls.
R-Value	Displays the R-value. R-value is a number or score used to quantitatively express the quality of speech in communications systems. Its used in digital networks that carry <i>Voice over IP</i> (VoIP) traffic. The R-value can range from 1 (worst) to 100 (best) and is based on the percentage of users who are satisfied with the quality of a test voice signal after it has passed through a network from a source (transmitter) to a destination (receiver). The R-value scoring method accurately portrays the effects of packet loss and delays in digital networks carrying voice signals.

9. The **Wireless Clients** table defines the following:

Client Count	Displays the number of wireless clients associated with the AP.
LT Clients	Displays the number of <i>Low Throughput</i> (LT) clients. This message can be delivered over a physical link, logical link or through a network node.
HT 20 Clients	Displays the number of clients using the HT 20 protection. This is a protection method to ensure backward compatibility with older 802.11abg radios. Using this protection, HT APs signal HT radios when to use protection mode.
Long Preamble Clients	Displays the number of clients using a long preamble. The radio preamble is a section of data in the packet header containing information the AP and client need when sending and receiving packets. A long preamble decreases throughput performance.
Long Slot Clients	Displays the number of clients having the long slot time. Slot time is the amount of time a device waits after a collision after retransmitting a packet. Long slot time decreases throughput.

Traffic

[Detailed Radio Statistics](#)

The Traffic screen provides the traffic utilization details of the radio. It also displays the different errors encountered during transmission and receiving of packets.

To view the traffic details:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Radios** from the left-hand side of the UI, and select **Radio Id**.
4. Select the **Traffic** sub-menu from the left navigation pane of the resulting screen.

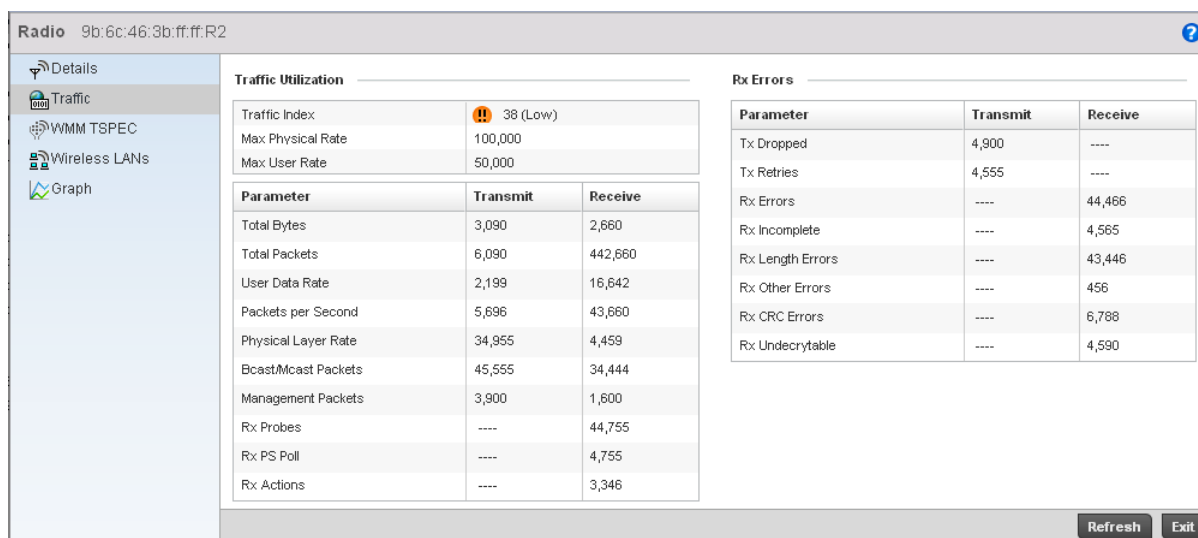


FIGURE 406 Access Point Radio Statistics Traffic screen

5. The **Traffic Utilization** area provides the following information:

Traffic Index	Displays the traffic utilization index, which measures how efficiently the traffic medium is used. It's defined as the percentage of current throughput relative to the maximum possible throughput. Traffic indices are: <ul style="list-style-type: none"> • 0–20 (very low utilization) • 20–40 (low utilization) • 40–60 (moderate utilization) • 60 and above (high utilization)
Max Physical Rate	Displays the maximum data rate at the physical layer.
Max User Rate	Displays the maximum permitted user data rate.
Total Bytes	Displays the total number of bytes processed by the AP radio.
Total Packets	Displays the total number of packets transmitted and received by the Access Point radio.
User Data Rate	Displays the average user data rate for both transmit and the receive operations.
Packets per Second	Displays the number of packets sent and received per second.
Physical Layer Rate	Displays data transfer rates at the physical layer for both transmit and receive operations.
Bcast/Mcast Packets	Displays the number of broadcast and multicast packets transmitted and received per second.
Management Packets	Displays the number of management packets processed by the radio per second.

Rx Probes	Displays the number of network probes received. A probe provides accurate statistics concerning a network's operation such as traffic analysis, top users, and protocol usage. Using this information, manage your network efficiently, ensuring peak operation and performance.
RX PS Poll	Displays the power save using the Power Save Poll mode. <i>Power Save Poll</i> (PSP) is a protocol, which helps to reduce the amount of time a radio needs to be powered. PSP allows the WiFi adapter to notify the AP when the radio is powered down. The access point holds any network packet sent to this radio.
Rx Actions	Displays the total number of action packets received.

6. The **Rx Errors** table defines the following:

Tx Dropped	Displays the number of transmitted packets dropped.
Tx Retries	Displays the average number of transmit retries. A high number indicates possible network or hardware problems
Rx Errors	Displays the number of errors received. The higher the error rate, the less reliable the connection or data transfer.
Rx Incomplete	Displays the number of incomplete packets received.
Rx Length Errors	Displays the number of length error packets received. Length errors are generated when the received frame length was less than or exceeded the Ethernet standard.
Rx Other Errors	Displays the number of additional errors that are not Rx Length Error or Rx CRC Error errors.
Rx CRC Errors	Displays the number of CRC errors received. Cyclic redundancy checks search for errors that has been transmitted on a communications link. A sending device applies a 16- or 32- bit polynomial to a block of transmitted data and appends the resulting cycling redundancy code to the block. The receiving end applies the same polynomial to the data and compares its result with the result appended by the sender. If those results agree, the data has been received successfully. If not, the sender can be notified to resend the block of data.
Rx Undecryptable	Displays the number of non-decryptable packets received.

WMM TSPEC

[Detailed Radio Statistics](#)

This *Traffic Specification* (TSPEC) screen displays the traffic specification parameters for the selected radio. The traffic specification allows an 802.11e client to signal its traffic requirements to the AP. It provides a set of parameters that define the characteristics of the traffic stream. The sender specifies parameters available for packet flow. Both the sender and the receiver use TSPEC.

The TSPEC screen provides the information about TSPEC counts and TSPEC types for the selected wireless client.

To view the TSPEC of the selected radio:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Radios** from the left-hand side of the UI, and select **Radio Id**.
4. Select the **WMM TSPEC** sub-menu from the left navigation pane of the resulting screen.

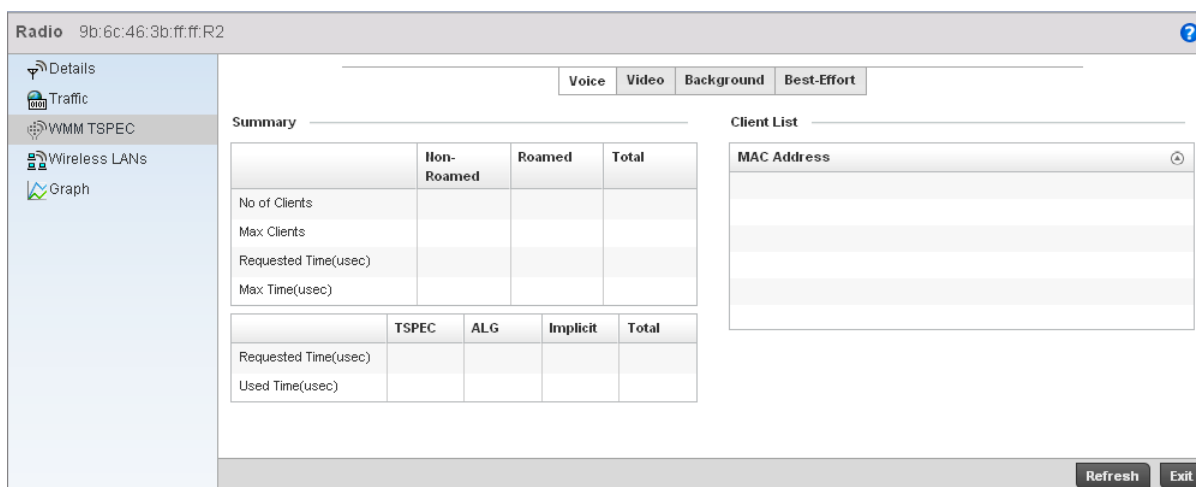


FIGURE 407 Access Point Radio Statistics TSPEC screen

5. The **Status** area provides the following:

- Voice** Displays the status of prioritization for voice traffic. A red 'X' indicates this feature is disabled. A green check mark indicates this feature is enabled.
- Video** Displays the status of prioritization for video traffic. A red 'X' indicates this feature is disabled. A green check mark indicates this feature is enabled.
- Best Effort** Displays the status of prioritization for best effort traffic. A red 'X' indicates this feature is disabled. A green check mark indicates this feature is enabled.
- Background** Displays the status of prioritization for background traffic. A red 'X' indicates this feature is disabled. A green check mark indicates this feature is enabled.

6. The **Utilization** table the following:

- Request Time TSPEC** Displays the time when the client asks the AP for its QoS requirements such as mean data rate, packet length, and extra allowance for retries through an *add traffic stream* (ADDTS) request when admission control is mandatory in the beacon. TSPEC allows a client to signal its traffic requirements to an AP. The AP decides whether the request is acceptable and transmits its decision to the client. The client can start high priority communication only when permitted by the AP.
- Used Time TSPEC** Displays the client's TSPEC usage. The client sends a *delete traffic stream* (DELTS) message when it's finished communicating.
- Request Time ALG** Displays the time when legacy clients or WMM clients, which do not support TSPEC, request QoS parameters from an existing application-level gateway (ALG) using a protocol like SIP.
- Used Time ALG** Displays the duration the legacy or WMM clients used a protocol with an existing ALG for obtaining QoS parameters (data rate, packet length and extra allowance of the AP).
- Request Time Implicit** Displays the time when an implicit request was generated. This is a feature to stimulate a TSPEC request for clients that do not support TSPEC and do not send traffic inspected by an ALG in an access controlled traffic class.
- Used Time Implicit** Displays the time taken by a wireless client to determine the QoS parameters of the AP using the Implicit TSPEC feature.

7. The **Roaming** table displays the following:

- Non-Roamed Clients** Displays the number of non-roaming clients associated with the access point radio. Non-roaming clients do not change their AP association from one AP to another. The non-roaming clients operate from a single location.
- Non-Roamed Time** Displays the time for which the non-roamed clients are in association with an AP radio.
- Roamed Clients** Displays the number of roaming clients associated with an AP radio. Roaming have been granted the roaming privilege. When roaming, it can toggle between normal and roaming modes. Roaming occurs when a client changes its AP association from one AP to another within the same WLAN.
- Roamed Time** Represents the time between the last *Real Time Transport Protocol (RTP)* packet seen on AP 1, and the first RTP packet seen on AP 2. It also includes the time the client takes to re-authenticate and re-associate with AP 2. (
- Max Non-Roamed Clients** Displays the maximum number of non roaming clients available for AP association.
- Max Non-Roamed Time** Displays the maximum time for which the non-roamed clients are in association with an AP radio.

Wireless LANs

Detailed Radio Statistics

The *Wireless LANs* screen provides WLAN details for the selected radio. It includes the WLAN name, its BSS ID and whether the WLAN is advertised.

To view WLAN statistics for the selected radio:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Radios** from the left-hand side of the UI, and select **Radio Id**.
4. Select the **Wireless LANs** sub-menu from the left navigation pane of the resulting screen.

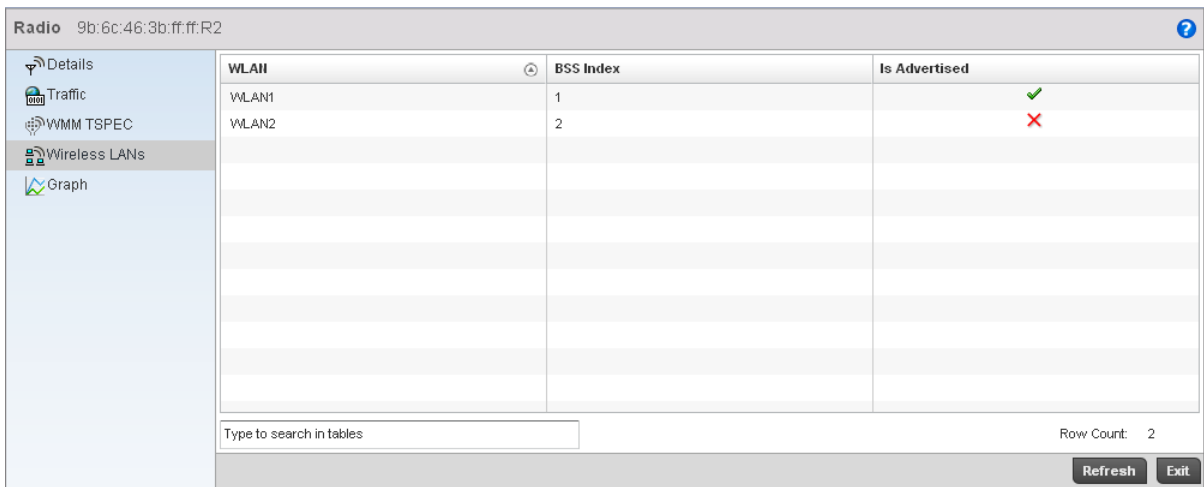


FIGURE 408 Access Point Radio Statistics Wireless LANs screen

5. The **Wireless LANs** screen displays the following:

WLAN	Displays the name of the WLAN the selected radio is using.
BSS Index	Displays the WLAN's <i>Basic Service Set</i> identifier (BSS). The BSS Index is the MAC address of the access point's radio servicing the BSS.
Is Advertised	Displays whether the WLAN is advertised through an SSID, which announces the availability of the wireless network to wireless clients.

AP Radio Graph

[Detailed Radio Statistics](#)

Use the Graph information to chart associated controller radio performance and diagnose radio performance issues. The graph uses a *Y-axis* and a *X-axis*. Select different parameters on the Y-axis and different polling intervals. To view a detailed graph for a client, select multiple parameters from the Y-axis and select a polling interval from the Y-axis.

To view the radio statistics graphically:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Radios** from the left-hand side of the UI, and select **Radio Id**.
4. Select the **Graph** sub-menu from the resulting screen.

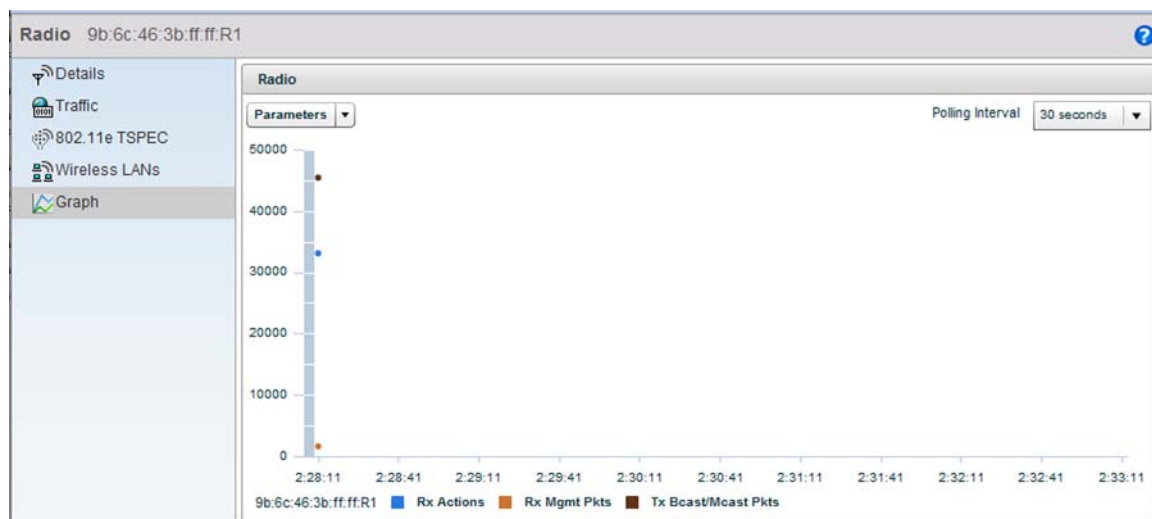


FIGURE 409 Access Point Radio Statistics Graph screen

Mesh

[Access Point Statistics](#)

The *Mesh* screen provides detailed statistics on each of the Mesh APs available. Use the following to review the performance of each AP interface.

To view the Mesh statistics:

1. Select the **Statistics** menu from the Web UI.

2. Select an **Access Point** node from the left navigation pane.
3. Select **Mesh** from the left-hand side of the UI.

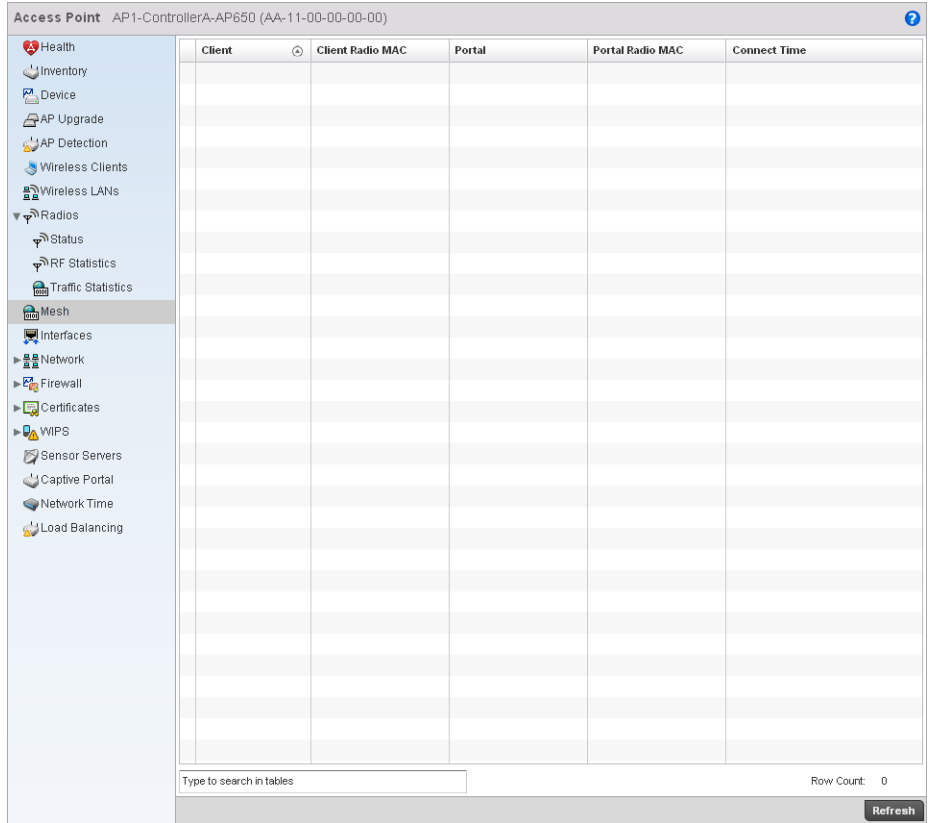


FIGURE 410 Access Point Mesh screen

4. The **Mesh** screen describes the following:

Interfaces

[Access Point Statistics](#)

The *Interface* screen provides detailed statistics on each of the interfaces available on an Access Point. Use the following to review the performance of each AP interface.

- [General Statistics](#)
- [Viewing Interface Statistics Graph](#)

General Statistics

Mesh

The *General* screen provides information on the interface such as its MAC address, type and TX/RX statistics.

To view the general interface statistics:

1. Select the **Statistics** menu from the Web UI.

2. Select an **Access Point** node from the left navigation pane.
3. Select **Interfaces** from the left-hand side of the UI.

The screenshot displays the configuration page for an interface named 'ge1' on an AP1-ControllerA-AP650. The interface is configured with the following parameters:

General	
Name	ge1
Interface MAC Address	33-11-00-11-00-11
IP Address	1.2.3.4
IP Address Type	
Secondary IPs	1 item
Hardware Type	ethernet
Index	1
Access VLAN	
Access Setting	Access
Administrative Status	true

Specification	
Media Type	one
Protocol	
MTU	1,000
Mode	Layer 2
Metric	
Maximum Speed	100M
Admin Speed	50
Operator Speed	true
Admin Duplex Setting	Half
Current Duplex Setting	Full

Traffic	
Good Octets Sent	200
Good Octets Received	13,200
Good Pkts Sent	200
Good Pkts Received	10,330
Mcast Pkts Sent	200
Mcast Pkts Received	200
Bcast Pkts Sent	200

Errors	
Bad Pkts Received	3,030
Collisions	200
Late Collisions	200
Excessive Collisions	200
Drop Events	100
Tx Undersize Pkts	4,030
Oversize Pkts	3,003
MAC Transmit Error	200
MAC Receive Error	200
Bad CRC	200

Receive Errors	
Rx Frame Errors	11
Rx Length Errors	12
Rx FIFO Errors	10
Rx Missed Errors	13
Rx Over Errors	14

Transmit Errors	
Tx Errors	18
Tx Dropped	17
Tx Aborted Errors	15
Tx Carrier Errors	16
Tx FIFO Errors	19
Tx Heartbeat Errors	20
Tx Window Errors	21

FIGURE 411 Access Point Interface screen

4. The **General** table describes the following:

Name	Displays the name of the interface.
Interface MAC Address	Displays the MAC address of the interface.
IP Address	IP address of the interface.
IP Address Type	Lists the IP address type of the interface
Hardware Type	Displays the hardware type.
Index	Displays the unique numerical identifier supporting the interface.
Access VLAN	Displays the interface the VLAN has access to.
Access Setting	Displays the mode of the VLAN as either <i>Access</i> or <i>Trunk</i> .
Administrative Status	Displays whether the interface is currently <i>UP</i> or <i>DOWN</i> .

5. The **Specification** table displays the following:

Media Type	Displays the physical connection type of the interface. Media types include: <i>Copper</i> - Used on RJ-45 Ethernet ports <i>Optical</i> - Used on fibre optic gigabit Ethernet ports
Protocol	Displays the name of the routing protocol adopted by the interface.
MTU	Displays the <i>maximum transmission unit</i> (MTU) setting configured on the interface. The MTU value represents the largest packet size that can be sent over a link. 10/100 Ethernet ports have a maximum setting of 1500.
Mode	The mode can be either: <i>Access</i> – This Ethernet interface accepts packets only from the native VLANs. <i>Trunk</i> – This Ethernet interface allows packets from a given list of VLANs that you can add to the trunk.
Metric	Displays the metric value associated with the route through this interface.
Maximum Speed	Displays the maximum speed at which the interface transmits or receives data.
Admin. Speed	Displays the speed setting used when using the administrative interface.
Operator Speed	Displays the current speed of the data transmitted and received over the interface.
Admin. Duplex Setting	Displays the administrator's duplex setting.
Current Duplex Setting	Displays the interface as either half duplex, full duplex, or unknown.

6. The **Traffic** table describes the following:

Good Octets Sent	Displays the number of octets (bytes) sent by the interface with no errors.
Good Octets Received	Displays the number of octets (bytes) received by the interface with no errors.
Good Pkts Sent	Describes the number of good packets transmitted.
Good Pkts Received	Describes the number of good packets received.
Mcast Pkts Sent	Displays the number of multicast packets sent through the interface.
Mcast Pkts Received	Displays the number of multicast packets received through the interface.
Bcast Pkts Sent	Displays the number of broadcast packets sent through the interface.
Bcast Pkts Received	Displays the number of broadcast packets received through the interface.
Packet Fragments	Displays the number of packet fragments transmitted or received through the interface.
Jabber Pkts	Displays the number of packets transmitted through the interface that is larger than the MTU through the interface.

7. The **Errors** table displays the following:

Bad Pkts Received	Displays the number of bad packets received through the interface.
Collisions	Displays the number of collisions.
Late Collisions	A late collision is any collision that occurs after the first 64 octets of data have been sent by the sending station. Late collisions are not normal, and are usually the result of out-of-specification cabling or a malfunctioning device.
Excessive Collisions	Displays the number of excessive collisions. Excessive collisions occur when the traffic load increases to the point that a single Ethernet network can not handle it efficiently.
Drop Events	Displays the number of dropped packets that are transmitted or received through the interface.
Tx Undersize Pkts	Displays the number of undersize packets transmitted through the interface.
Oversize Pkts	Displays the number of oversize packets.
MAC Transmit Error	Displays the number of failed transmits due to an internal MAC sublayer error (not a late collision, excessive collisions or carrier sense error).
MAC Receive Error	Displays the number of failed received packets due to an internal MAC sublayer (not a late collision, excessive collisions or carrier sense error).
Bad CRC	Displays the CRC error. The <i>Cyclical Redundancy Check (CRC)</i> is the 4 byte field at the end of every frame. The receiving station uses it to interpret if the frame is valid. If the CRC value computed by the interface does not match the value at the end of the frame, it's considered a bad CRC.

8. The **Receive Errors** table displays the following:

Rx Frame Errors	Displays the number of frame errors received at the interface. A frame error occurs when a byte of data is received in unexpected format.
Rx Length Errors	Displays the number of length errors received at the interface. Length errors are generated when the received frame length was less than or exceeded the Ethernet standard.
Rx FIFO Errors	Displays the number of FIFO errors received at the interface. <i>First-in First-Out</i> queueing is an algorithm that involves buffering and forwarding of packets in the order of arrival. FIFO entails no priority for traffic. There is only one queue, and all packets are treated equally.
Rx Missed Errors	Displays the number of missed packets. Packets are missed when the hardware received FIFO has insufficient space to store the incoming packet.
Rx Over Errors	Displays the number of overflow errors. An overflow occurs when packet size exceeds the allocated buffer size.

9. The **Transmit Errors** table displays the following:

Tx Errors	Displays the number of packets with errors transmitted on the interface.
Tx Dropped	Displays the number of transmitted packets dropped from the interface.
Tx Aborted Errors	Displays the number of packets aborted on the interface because a clear-to-send request was not detected.
Tx Carrier Errors	Displays the number of carrier errors on the interface. This generally indicates bad Ethernet hardware or cabling.

- Tx FIFO Errors** Displays the number of FIFO errors received at the interface. FIFO is an algorithm that involves buffering and forwarding packets in the order of arrival. FIFO provides no priority for traffic. There is only one queue, and all packets are treated equally.
- Tx Heartbeat Errors** Displays the number of heartbeat errors. This generally indicates a software crash or packets stuck in an endless loop.
- Tx Window Errors** Displays the number of window errors transmitted. TCP uses a sliding window flow control protocol. In each TCP segment, the receiver specifies the amount of additional received data (in bytes) in the receive window field the receiver is willing to buffer for the connection. The sending host can only send up to that amount. If the sending host transmits more data before receiving an acknowledgement from the receiving host, it constitutes a window error.

Viewing Interface Statistics Graph

Mesh

The **Network Graph** tab displays interface statistics graphically. To view a detailed graph for an interface, select an interface and drop it on to the graph. The graph has *Port Statistics* as the Y-axis and the *Polling Interval* as the X-axis. Select different parameters on the Y-axis and different polling intervals as needed.

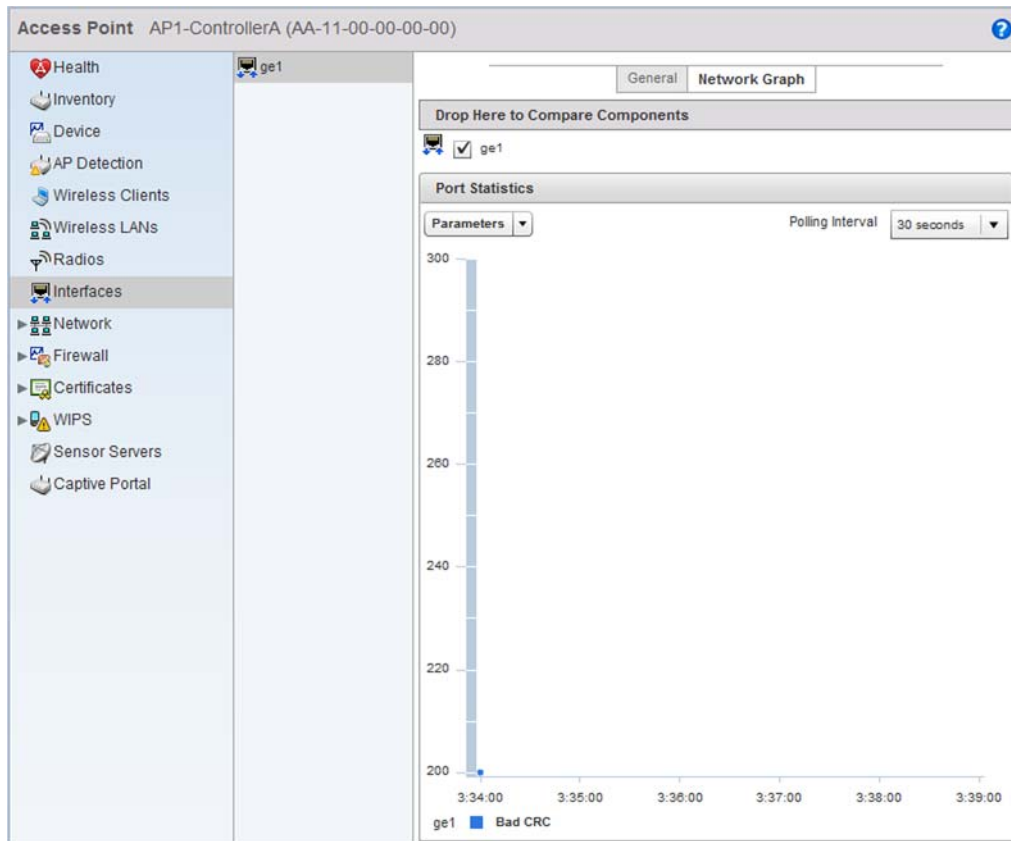


FIGURE 412 Access Point Interface Graph screen

Network

[Access Point Statistics](#)

Use the *Network* screen to view information for *ARP*, *Route Entry*, *Bridge*, *DHCP*, *CDP* and *LLDP*. Each of these screens provide enough statistics to troubleshoot issues related to the following four features:

- [ARP Entries](#)
- [Route Entries](#)
- [Bridge](#)
- [DHCP Options](#)
- [Cisco Discovery Protocol](#)
- [Link Layer Discovery Protocol](#)

ARP Entries

[Network](#)

ARP is a networking protocol for determining a network host's hardware address when its IP address or network layer address is known. The ARP screen displays entries configured for this wireless controller.

To view the ARP statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Network > ARP Entries** from the left-hand side of the UI.

- Issues IP addresses
- Throttles bandwidth
- Permits access to other networks
- Times out old logins

The Bridging screen also provides information about the *Multicast Router* (MRouter), which is a router program that distinguishes between multicast and unicast packets and how they should be distributed along the Multicast Internet. Using an appropriate algorithm, a multicast router instructs a switching device what to do with the multicast packet.

This screen is partitioned into the following:

- [Details](#)
- [MAC Address](#)

Details

To view the Bridge details:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Network > Bridge** from the left-hand side of the UI, and select the **Details** tab.

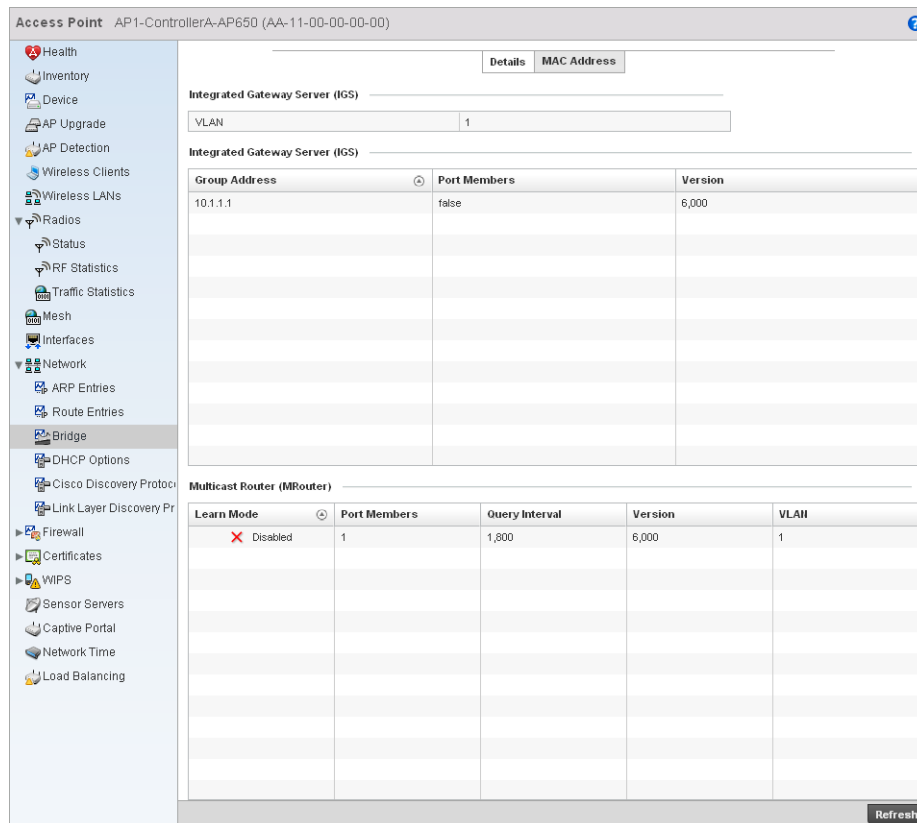


FIGURE 415 Access Point Network Bridge Details screen

4. The **Integrated Gateway Server (IGS)** table displays the following:

VLAN	Displays the VLAN where the multicast transmission is conducted.
Group Address	Displays the Multicast Group ID supporting the statistics displayed. This group ID is the multicast address hosts are listening to.
Port Members	Displays the ports on which multicast clients have been discovered by the wireless controller. Displays the interface name. For example, ge1, radio 1, etc.
Version	Displays the IGMP version in use.

5. The **Multicast Router (MRouter)** table describes the following:

Learn Mode	Displays the learning mode used by the router. Either <i>Static</i> or <i>PIM-DVMRP</i> .
Port Members	Displays the ports on which multicast clients have been discovered by the wireless controller. Displays the interface name. For example, ge1 or radio 1.
Query Interval	Displays the periodic IGMP query interval value.
Version	Displays the IGMP version in use.
VLAN	Displays the VLAN on which the multicast transmission is being done.

MAC Address

To view the Bridge MAC address details:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Network > Bridge** from the left-hand side of the UI, and select the **MAC Address** tab.

Bridge Name	MAC Address	Interface	VLAN	Forwarding
bridge1	11-22-33-44-55-66	eth0	1	true

FIGURE 416 Access Point Network Bridge MAC Address screen

4. The **MAC Address** tab displays the following:

Bridge Name	Displays the name of the network bridge.
MAC Address	Displays the MAC address of the bridge selected.
Interface	Displays the interface where the bridge transferred packets.
VLAN	Displays the VLAN the bridge belongs to.
Forwarding	Displays whether the bridge is forwarding packets. A bridge can only forward packets.

DHCP Options

Network

The controller contains an internal *Dynamic Host Configuration Protocol* (DHCP) server. The DHCP server can provide the dynamic assignment of IP addresses automatically. This is a protocol that includes IP address allocation and delivery of host-specific configuration parameters from a DHCP server to a host. Some of these parameters are IP address, gateway and network mask.

The **DHCP Options** screen provides the DHCP server name, image file on the DHCP server, and its configuration.

To view a network's DHCP Options:

1. Select the **Statistics** menu from the Web UI.

2. Select an **Access Point** node from the left navigation pane.
3. Select **Network > DHCP Options**.

Server Information	Image File	Configuration	Legacy Adoption	Adoption
server_information_1	image_1	configuration_1		
server_information_2	image_2	configuration_2		

FIGURE 417 Access Point Network DHCP Options screen

4. The **DHCP Options** screen displays the following:

Server Information Displays the IP address of the DHCP server.

Image File Displays the image file name. BOOTP or the bootstrap protocol can be used to boot diskless clients. An image file is sent from the boot server. The image file contains the image of the operating system the client will run. DHCP servers can be configured to support BOOTP.

Configuration Displays the name of the configuration file on the DHCP server.

Cluster Configuration Displays the name of the cluster configuration file on the DHCP server if the server is a part of a cluster.

Cisco Discovery Protocol

Network

To view a network's CDP Statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.

3. Select **Network > Cisco Discovery Protocol**.

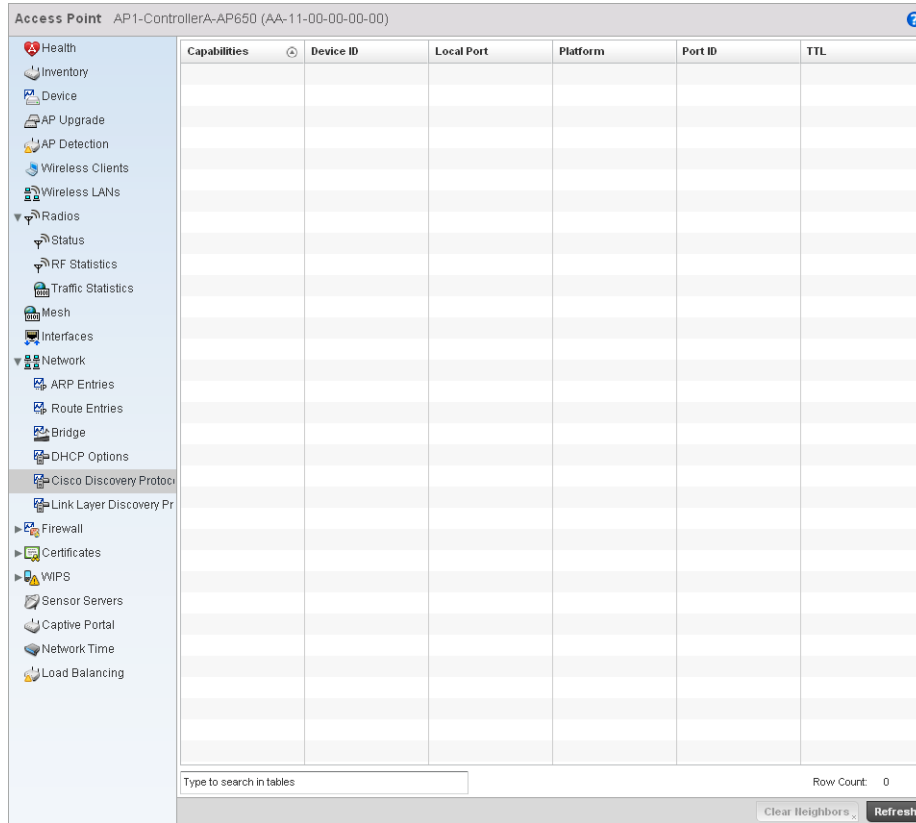


FIGURE 418 Access Point Network Cisco Discovery Protocol screen

4. The **Cisco Discovery Protocol** screen displays the following:

- Capabilities** Displays the capabilities code for the device either Router, Trans Bridge, Source Route Bridge, Switch, Host, IGMP or Repeater.
- Device ID** Displays the configured device ID or name for each device in the table.
- Local Port** Displays the local port name for each CDP capable device.
- Platform** Displays the model number of the CDP capable device.
- Port ID** Displays the identifier for the local port.
- TTL** Displays the time to live for each CDP connection.
- Clear Neighbors** Click Clear Neighbors to remove all known CDP neighbors from the table.

Link Layer Discovery Protocol

Network

To view a network’s Link Layer Discovery Protocol statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.

3. Select **Network > Link Layer Discovery Protocol**.

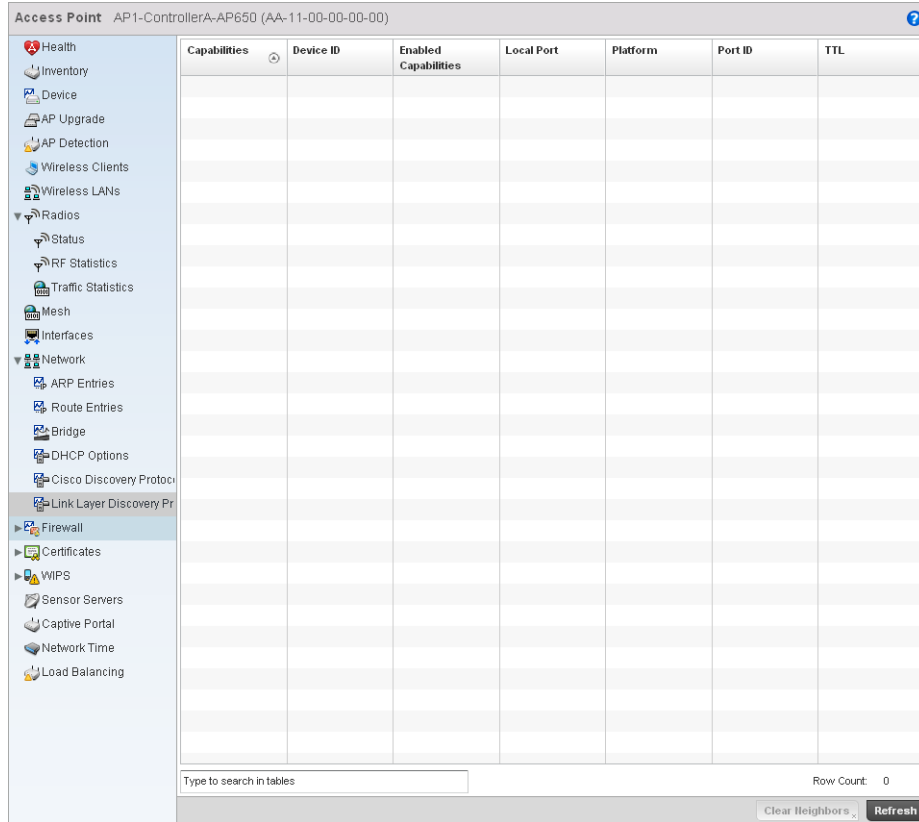


FIGURE 419 Access Point Link Layer Discovery screen

4. The **Link Layer Discovery Protocol** screen displays the following:

- Capabilities** Displays the capabilities code for the device either Router, Trans Bridge, Source Route Bridge, Switch, Host, IGMP or Repeater.
- Device ID** Displays the configured device ID or name for each device in the table.
- Enabled Capabilities** Displays which of the device capabilities are currently enabled.
- Local Port** Displays the local port name for each LLDP capable device.
- Platform** Displays the model number of the LLDP capable device.
- Port ID** Displays the identifier for the local port.
- TTL** Displays the time to live for each LLDP connection.
- Clear Neighbors** Click Clear Neighbors to remove all known LLDP neighbors from the table.

Firewall

[Access Point Statistics](#)

A firewall blocks unauthorized access while permitting authorized communications. It's a device or set of devices configured to permit or deny computer applications based on a set of rules.

This screen is partitioned into the following:

- [Packet Flows](#)
- [Denial of Service](#)
- [IP Firewall Rules](#)
- [MAC Firewall Rules](#)
- [NAT Translations](#)
- [DHCP Snooping](#)

Packet Flows

Firewall

The *Packet Flows* screen displays a bar graph for the different packet types flowed through the Access Point. Use this information to assess the traffic patterns supported by the Access Point.

The *Total Active Flows* graph displays the total number of flows supported. Other bar graphs display for each individual packet type.

To view the packet flows statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Firewall > Packet Flows**.

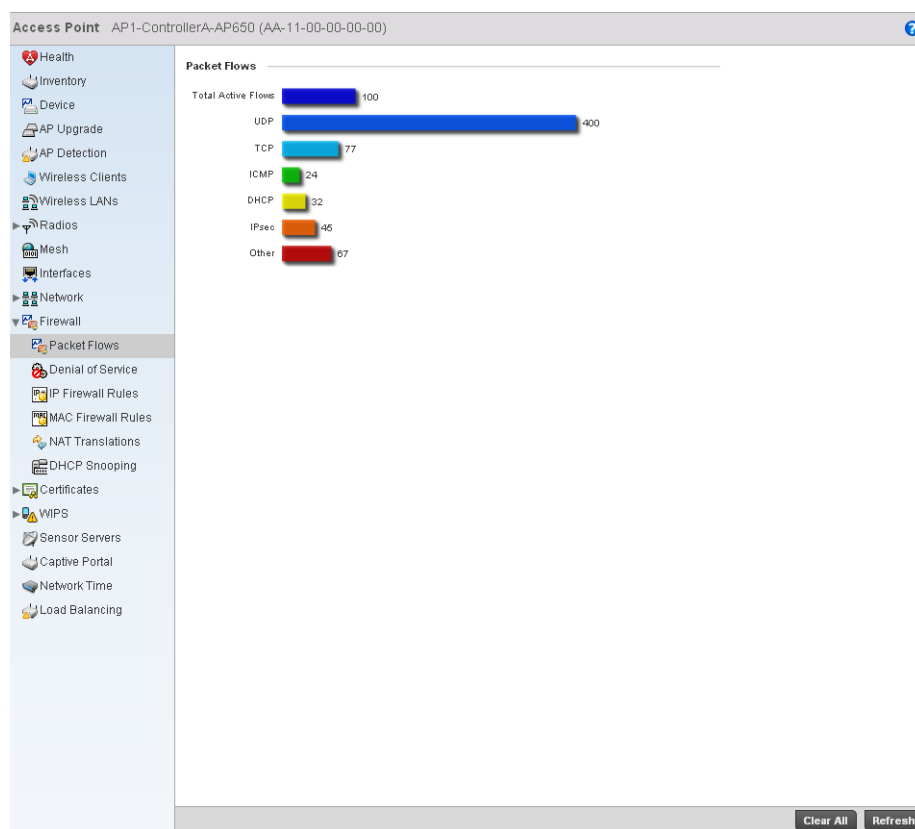


FIGURE 420 Access Point Firewall Statistics Packet Flow screen

Denial of Service

Firewall

A *denial-of-service attack* (DoS attack) or distributed denial-of-service attack is an attempt to make a computer resource unavailable to its intended users. Although the means to carry out a DoS attack may vary, it generally consists of concerted efforts to prevent an Internet site or service from functioning efficiently.

One common method involves saturating the target's machine with external communications requests, so it cannot respond to legitimate traffic or responds so slowly as to be rendered effectively unavailable. DoS attacks are implemented by either forcing the targeted computer(s) to reset, or consume its resources so it can't provide its intended service.

The DoS screen displays the types of attack, number of times it occurred and the time of last occurrence.




FIGURE 421 Access Point Firewall Denial of Service screen

The **Denial of Service** screen displays the following:

- Attack Type** Displays the *Denial of Service* (DoS) attack type.
- Count** Displays the number of times the firewall has observed each DoS attack.
- Last Occurrence** Displays the amount of time since the DoS attack has been observed by the firewall.

IP Firewall Rules

Firewall

Create firewall rules to permit a computer to send traffic to, or receive traffic from, programs, system services, computers or users. Firewall rules can be created to take one of the three actions listed below that match the rule's criteria:

- Allow a connection
- Allow a connection only if it is secured through the use of Internet Protocol security
- Block a connection

Rules can be created for either inbound or outbound traffic.

To view the IP firewall rules:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Firewall > IP Firewall Rules** from the left-hand side of the UI.

Precedence	Friendly String	Hit Count
	firewall1	10

Type to search in tables Row Count: 1

[Refresh](#)

FIGURE 422 Access Point IP Firewall Rules screen

4. The **IP Firewall Rules** screen displays the following:

Precedence	Displays the precedence applied to packets. The rules within an <i>Access Control Entries</i> (ACL) are based on precedence. Every rule has a unique precedence value between 1 and 5000. You cannot add two rules with the same precedence.
Friendly String	This is a string that provides more information as to the contents of the rule.
Hit Count	Displays the number of times each WLAN ACL has been triggered.

MAC Firewall Rules

Firewall

The ability to allow or deny a system by its MAC address ensures malicious or unwanted users are unable to bypass security filters. Firewall rules can be created to support one of the three actions listed below that match the rule's criteria:

- Allow a connection
- Allow a connection only if it is secured through the MAC firewall security
- Block a connection

To view the MAC Firewall Rules:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Firewall > MAC Firewall Rules** from the left-hand side of the UI.

Precedence	Friendly String	Hit Count
	firewall1	10

Access Point AP1-ControllerA-AP650 (AA-11-00-00-00-00)

macaci5

Type to search in tables

Row Count: 1

Refresh

FIGURE 423 Access Point MAC Firewall Rules screen

4. The **MAC Firewall Rules** screen provides the following information:

- Precedence** Displays the precedence value, which are applied to packets. The rules within an *Access Control Entries (ACL)* list are based on their precedence values. Every rule has a unique precedence value between 1 and 5000. You cannot add two rules with the same precedence.
- Friendly String** Displays a string providing additional information on rule contents.
- Hit Count** Displays the number of times each WLAN ACL has been triggered.

NAT Translations

Firewall

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Firewall > NAT Translations**.

Protocol	Forward Source IP	Forward Source Port	Forward Dest IP	Forward Dest Port	Reverse Source IP	Reverse Source Port	Reverse Dest IP	Reverse Dest Port
UDP	7.7.7.7	777	3.3.3.3	333	8.8.8.8	888	4.4.4.4	444

FIGURE 424 Access Point Firewall NAT Translation screen

4. The **NAT Translations** screen displays the following:

Protocol	Displays the IP protocol type, either UDP or TCP.
Forward Source IP	Displays the internal network IP address for forward facing NAT translations in the Forward Source IP column.
Forward Source Port	Displays the internal network port for forward facing NAT translations in the Forward Source Port column.
Forward Dest IP	Displays the external network destination IP address for forward facing NAT translations in the Forward Dest IP column.
Forward Dest Port	Displays the external network destination port for forward facing NAT translations in the Forward Dest Port column.
Reverse Source IP	Displays the internal network IP address for reverse facing NAT translations in the Reverse Source IP column.
Reverse Source Port	Displays the internal network port for reverse facing NAT translations in the Reverse Source Port column.
Reverse Dest IP	Displays the external network destination IP address for reverse facing NAT translations in the Reverse Dest IP column.
Reverse Dest Port	Displays the external network destination port for reverse facing NAT translations in the Reverse Dest Port column.

DHCP Snooping

Firewall

When DHCP servers are allocating IP addresses to clients on the LAN, DHCP snooping can be configured on LAN controllers to allow only clients with specific IP or MAC addresses.

MAC Address	Node Type	IP Address	Netmask	VLAN	Lease Time	Time Elapsed Since Last Update
22-33-44-55-66-77	Router	4.4.4.4	0	100	5m 0s	30s

FIGURE 425 Access Point Firewall DHCP Snooping screen

The **DHCP Snooping** screen displays the following:

MAC Address	Displays the MAC address of the client.
Node Type	Displays the NetBios node with the IP pool from which IP addresses can be issued to client requests on this interface.
IP Address	Displays the IP address used for DHCP discovery, and requests between the DHCP server and DHCP clients.
Netmask	Displays the subnet mask used for DHCP discovery, and requests between the DHCP server and DHCP clients.
VLAN	Displays the controller interface used for the newly created DHCP configuration.
Lease Time	When a DHCP server allocates an address for a DHCP client, the client is assigned a lease (which expires after a designated interval defined by the administrator). The lease time is the time an IP address is reserved for re-connection after its last use. Using very short leases, DHCP can dynamically reconfigure networks in which there are more computers than there are available IP addresses. This is useful, for example, in education and customer environments where client users change frequently. Use longer leases if there are fewer users.
Time Elapsed Since Last Update	Displays the time the server was last updated.

Certificates

Access Point Statistics

The *Secure Socket Layer* (SSL) secures transactions between Web servers and browsers. SSL uses a third-party certificate authority to identify one (or both) ends of a transaction. A browser checks the server issued certificate before establishing a connection.

This screen is partitioned into the following:

- [Trustpoints](#)
- [RSA Keys](#)

Trustpoints

Certificates

Each certificate is digitally signed by a trustpoint. The trustpoint signing the certificate can be a certificate authority, corporation or individual. A trustpoint represents a CA/identity pair containing the identity of the CA, CA-specific configuration parameters and an association with an enrolled identity certificate.

The screenshot shows the 'Access Point' configuration interface for 'AP1-ControllerA-AP650 (AA-11-00-00-00)'. The left sidebar contains a navigation tree with 'Trustpoints' highlighted. The main content area is titled 'Certificate Details' and shows the following information:

Certificate Details	
Subject Name	CN=US, ST=California, L=San Jose, O=Motorola Incorporated, OU=https://www.motorola.com,
Alternate Subject Name	M1
Issuer Name	Motorola
Serial Number	dd111
RSA Key Used	false
Is CA	✓ True
Is Self Signed	✓ True
Server Certificate Present	✓ True
CRL Present	✓ True
Validity	
Valid From	2007
Valid Until	2011
Certificate Authority (CA) Details	
Subject Name	CN=US, ST=California, L=San Jose, O=Motorola Incorporated, OU=https://www.motorola.com,
Alternate Subject Name	M3
Issuer Name	Motorola
Serial Number	1,244
Certificate Authority Validity	
Valid From	2009
Valid Until	2011

A 'Refresh' button is located at the bottom right of the screen.

FIGURE 426 Access Point Certificate Trustpoint screen

The **Certificate Details** field displays the following:

Subject Name	Lists details about the entity to which the certificate is issued.
Alternate Subject Name	Displays alternative details to the information specified under the Subject Name field.
Issuer Name	Displays the name of the organization issuing the certificate.
Serial Number	The unique serial number of the certificate issued.
RSA Key Used	Displays the name of the key pair generated separately, or automatically when selecting a certificate.
IS CA	States whether this certificate is a authority certificate.
Is Self Signed	States whether the certificate is self-signed. True indicates the certificate is self-signed.
Server Certificate Present	Displays if the server certificate is present. True indicates the certificate is present.
CRL Present	Displays whether this functionality is present or not. The <i>Certificate Revocation List</i> (CRL) uses a public key infrastructure to maintain access to network servers.

RSA Keys

Certificates

Rivest, Shamir, and Adleman (RSA) is an algorithm for public key cryptography. It's the first algorithm known to be suitable for signing, as well as encryption.

The *RSA Keys* screen displays a list of RSA keys installed in the selected wireless controller. RSA Keys are generally used for establishing a SSH session, and are a part of the certificate set used by RADIUS, VPN and HTTPS.

To view the RSA Key details:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Certificates > RSA Keys** from the left-hand side of the UI.

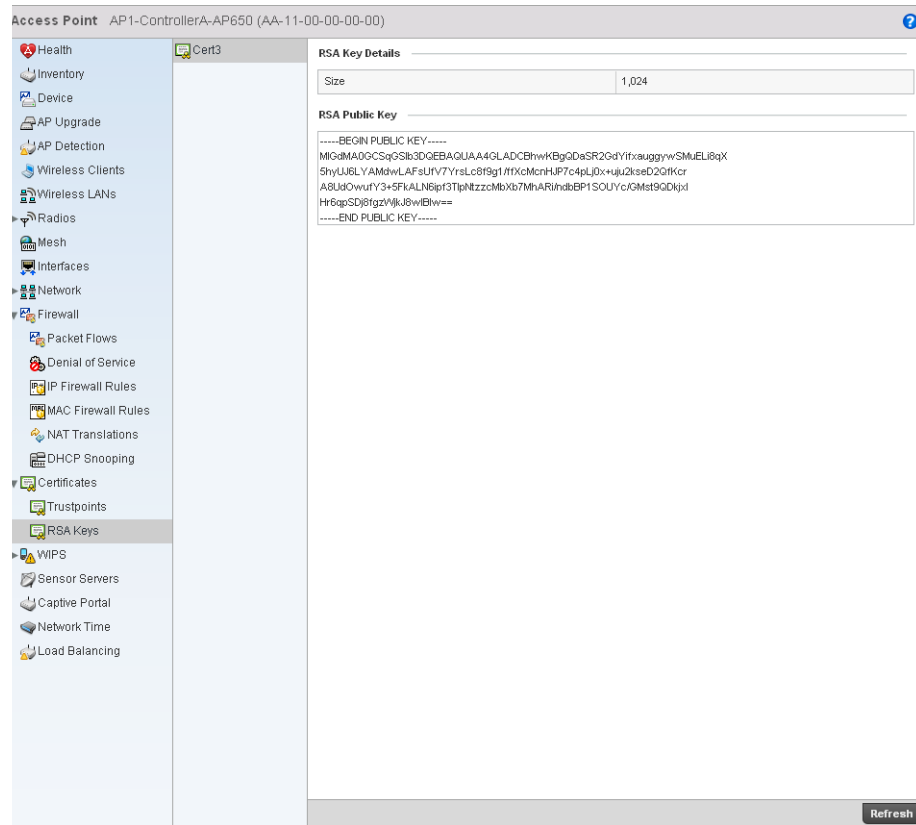


FIGURE 427 Access Point Certificates RSA Key screen

- The **RSA Key Details** table displays the size (in bits) of the desired key. If not specified, a default key size of 1024 is used.

The *RSA Public Key* field lists the public key used for encrypting messages.

WIPS

[Access Point Statistics](#)

A *Wireless Intrusion Prevention System (WIPS)* monitors the radio spectrum for the presence of unauthorized Access Points and take measures to prevent an intrusion. Unauthorized attempts to access the WLAN is generally accompanied by anomalous behavior as intruding clients try to find network vulnerabilities. When the parameters exceed a configurable threshold, the controller generates a SNMP trap and reports the results via management interfaces. Basic WIPS functionality does not require monitoring APs, and does not perform off-channel scanning.

The WIPS screen provides details about the blacklisted clients (unauthorized access points) intruded into the network. The details include the name of the blacklisted client, the time when the client was blacklisted, the total time the client remained in the network, etc. The screen also provides WIPS event details.

The WIPS screen is partitioned into:

- [Client Blacklist](#)
- [WIPS Events](#)

Client Blacklist

WIPS

The **Client Blacklist** screen displays blacklisted client data. It includes the name of the client, time when the blacklist event occurred and the duration the blacklisted client remained in the network.

To view the Client Blacklist screen:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **WIPS > Client Blacklist** from the left-hand side of the UI.

Event Name	Blacklisted Client	Time Blacklisted	Total Time	Time Left
dos-eapol-start-storm	44-55-44-55-44-55	Thu Jun 10 2010 12:26:28 PM	2h 0m 0s	1h 0m 0s
null-probe-response	44-55-44-55-44-55	Thu Jun 10 2010 12:26:28 PM	40m 0s	20m 0s

FIGURE 428 Access Point WIPS Client Blacklist screen

4. The **Client Blacklist** screen provides the following information:

Event Name	Displays the name of the wireless intrusion event.
Blacklisted Client	Displays the MAC address of the intruding unauthorized access point.
Time Blacklisted	Displays the time when this client was blacklisted.
Total Time	Displays the total time the unsanctioned AP remained in the WLAN.
Time Left	Displays the remaining blacklist duration time. After this time elapses, the client is removed from the blacklist.

WIPS Events

WIPS

The WIPS Events screen details the wireless intrusion.

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **WIPS > WIPS Events** from the left-hand side of the UI.

Event Name	Reporting AP	Originating Device	Detector Radio	Time Reported
dos-eapol-start-storm	API-ControllerA-AP650	33-44-33-44-33-44	1	Thu Jun 10 2010 12:26:28 PM
null-probe-response	API-ControllerA-AP650	33-44-33-44-33-44	1	Thu Jun 10 2010 12:26:28 PM

FIGURE 429 Access Point WIPS Events screen

4. The WIPS Events screen provides the following:

- Event Name** Displays the name of the wireless intrusion event.
- Reporting AP** Displays the MAC address of the AP reporting this intrusion.
- Originating Device** Displays the MAC address of the intruding device.
- Time Reported** Displays the time when the intrusion was detected.

Sensor Servers

Access Point Statistics

Sensor servers allow an administrator to monitor and download data from multiple sensors and remote locations using Ethernet TCP/IP or serial communications. Repeaters are available to extend transmission range and combine sensors with various frequencies on the same receiver.

To view the sensor server statistics of an AP:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Sensor Servers** from the left-hand side of the UI.

IP Address	Port	Status
1.1.1.1	8,443	connected
2.2.2.2	443	not connected

FIGURE 430 Access Point Sensor Servers screen

4. The **Sensor Servers** screen displays the following:

IP Address	Displays the IP address of the sensor server.
Port	Displays the port on which this server is listening.
Status	Displays whether the server is <i>UP</i> or <i>DOWN</i> .

Captive Portal

[Access Point Statistics](#)

A captive portal forces a HTTP client to use a special Web page for authentication before using the Internet. A captive portal turns a Web browser into a client authenticator. This is done by intercepting packets regardless of the address or port, until the user opens a browser and tries to access the Internet. At that time, the browser is redirected to a Web page.

To view the captive portal statistics of an access point:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Captive Portal** from the left-hand side of the UI.

Access Point AP1-ControllerA-AP650 (AA-11-00-00-00)

Client MAC	Client IP	Captive Portal	Authentication	WLAN	VLAN	Remaining Time
AA-11-11-00-00-00	1.1.1.1	default	Success	WLAN3	1	1m 40s
AA-11-12-00-00-00	1.1.1.1	default	Pending	WLAN4	2	3m 20s

Type to search in tables Row Count: 2

Refresh

FIGURE 431 Access Point Captive Portal screen

4. The **Captive Portal** screen supports the following:

Client MAC	Displays the client's MAC address.
Client IP	Displays the client's IP address.
Captive Portal	Displays the IP address of the captive portal page.
Authentication	Displays the authentication status of the wireless client.
WLAN	Displays the name of the requesting client's WLAN.
VLAN	Displays the name of the requesting client's VLAN.
Remaining Time	Displays the time after which the client is disconnected from the Internet.

Network Time

System Statistics

Network Time Protocol (NTP) is central to networks that rely on their wireless controller to supply system time. Without NTP, controller time is unpredictable, which can result in data loss, failed processes, and compromised security. With network speed, memory, and capability increasing at an exponential rate, the accuracy, precision, and synchronization of network time is essential in a controller-managed enterprise network. The wireless controller can use a dedicated server to supply system time. The controller can also use several forms of NTP messaging to sync system time with authenticated network traffic.

Viewing NTP Status

Network Time

The NTP Status screen displays performance (status) information relative to the AP's current NTP association. Verify the controller's NTP status to assess the controller's current NTP resource.

To view the NTP status of a managed network:

- 1. Select the **Statistics** menu from the Web UI.
- 2. Select an **Access Point** node from the left navigation pane.
- 3. Select **Network Time > NTP Status** from the left-hand side of the UI.

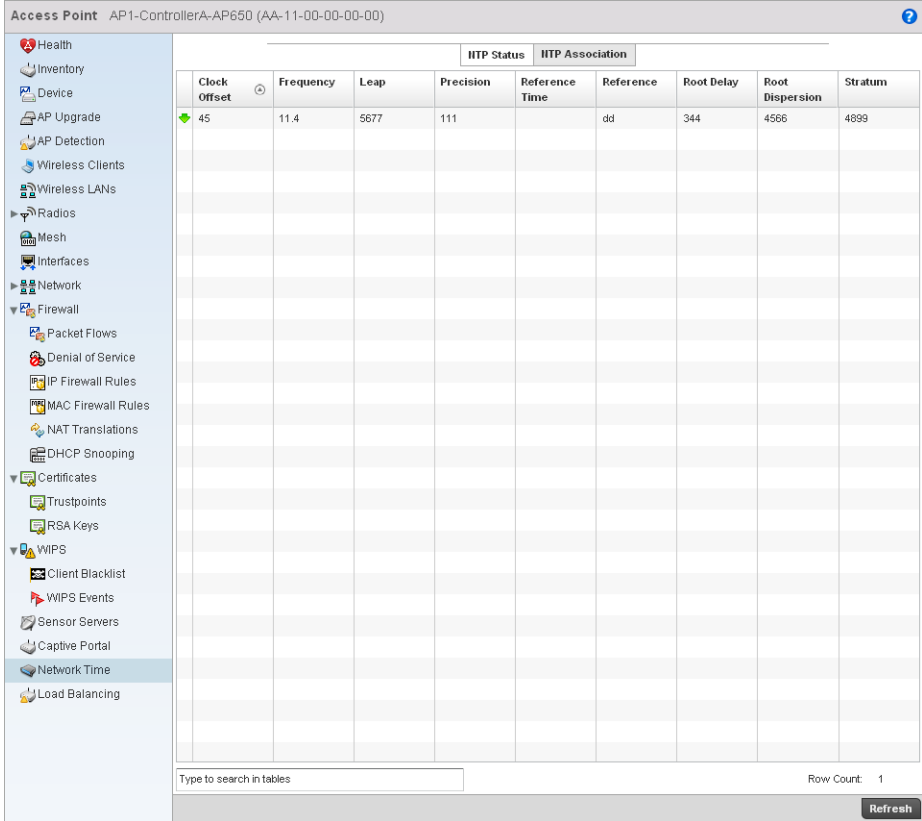


FIGURE 432 System NTP screen

4. Refer to the **NTP Status** table to review the accuracy and performance of the controller's synchronization with an NTP server.

Clock Offset	Displays the time differential between the controller time and the NTP resource.
Frequency	An SNTP server clock's skew (difference) for the controller.
Leap	Indicates if a second is added or subtracted to SNTP packet transmissions, or if transmissions are synchronized.
Precision	Displays the precision of the controller's time clock (in Hz). The values that normally appear in this field range from -6 for mains-frequency clocks to -20 for microsecond clocks.
Reference Time	Displays the time stamp the local clock was last set or corrected.
Reference	Displays the address of the time source the controller is synchronized to.
Root Delay	The total round-trip delay in seconds. This variable can take on both positive and negative values, depending on relative time and frequency offsets. The values that normally appear in this field range from negative values (a few milliseconds) to positive values (several hundred milliseconds).
Root Dispersion	The difference between the time on the root NTP server and its reference clock. The reference clock is the clock used by the NTP server to set its own clock.
Status Stratum	Displays how many hops the controller is from its current NTP time source.

Viewing NTP Associations

Network Time

The interaction between the controller and an SNTP server constitutes an association. SNTP associations can be either peer associations (the controller synchronizes to another system or allows another system to synchronize to it), or a server associations (only the controller synchronizes to the SNTP resource, not the other way around).

To view the NTP associations:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Network > NTP Association** from the left-hand side of the UI.

Access Point AP1-ControllerA-AP650 (AA-11-00-00-00)										
					NTP Status		NTP Association			
Delay Time	Display	Offset	Poll	Reach	Reference IP Address	Server IP Address	State	Status	Time	
10	45	67	44	445	12.34.44.44	12.2.2.2	455	ss	now	

FIGURE 433 AP Network Time screen

4. The **NTP Associations** screen provides the controller's current NTP associations.

This screen provides the following:

Delay Time	Displays the round-trip delay (in seconds) for SNTP broadcasts between the SNTP server and the wireless controller.
Dispersion	Displays the time difference between the peer NTP server and the onboard wireless controller clock.
Offset	Displays the calculated offset between the wireless controller and the SNTP server. The controller adjusts its clock to match the server's time value. The offset gravitates towards zero overtime, but never completely reduces its offset to zero.
Poll	Displays the maximum interval between successive messages (in seconds) to the nearest power of two.
Reach	Displays the status of the last eight SNTP messages. If an SNTP packet is lost, the lost packet is tracked over the next eight SNTP messages.
Reference IP Address	Displays the address of the time source the wireless controller is synchronized to.
Server IP Address	Displays the numerical IP address of the SNTP resource (server) providing SNTP updates to the wireless controller.

State

Displays the NTP association status. The state can be one of the following:

- *Synced* – Indicates the wireless controller is synchronized to this NTP server.
- *Unsynced* – Indicates the wireless controller has chosen this master for synchronization. However, the master itself is not yet synchronized to UTC.
- *Selected* – Indicates this NTP master server will be considered the next time the wireless controller chooses a master to synchronize with.
- *Candidate* – Indicates this NTP master server may be considered for selection the next time the wireless controller chooses a NTP master server.
- *Configured* – Indicates this NTP server is a configured server.

Stratum

Displays the NTP peer's stratum level.

When

Displays the timestamp of the last NTP packet received from the NTP peer.

Load Balancing

Access Point Statistics

To view the **Load Balancing** statistics of an access point:

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Load Balancing** from the left-hand side of the UI.

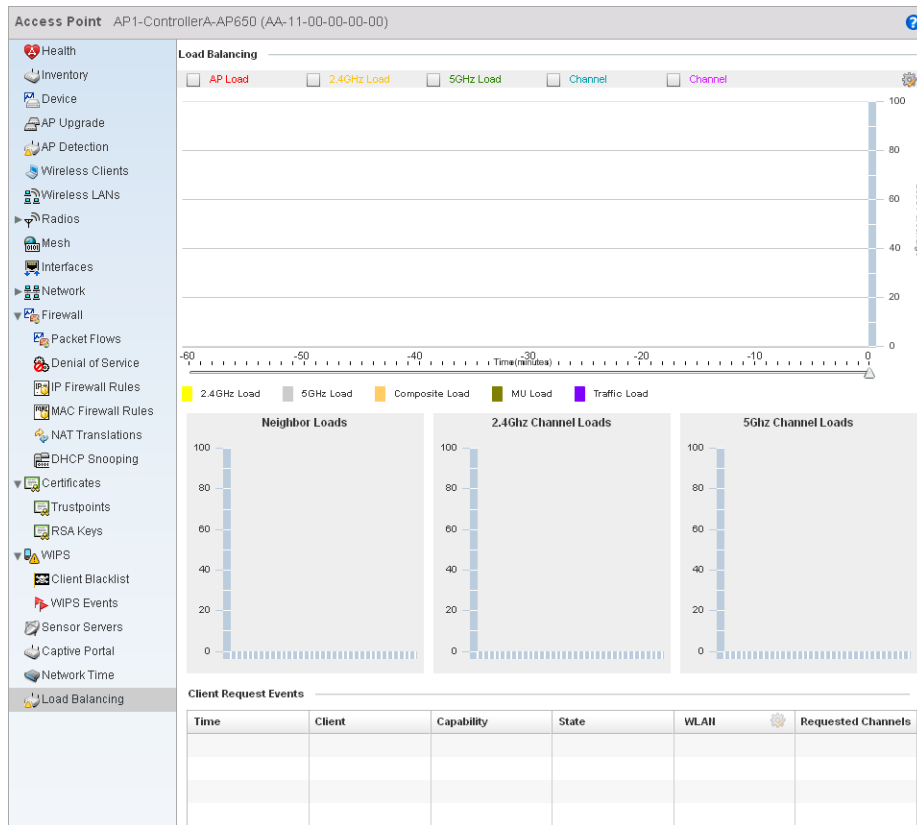


FIGURE 434 Access Point Load Balancing screen

4. The **Load Balancing** screen supports the following:

Load Balancing	In the Load Balancing section check the boxes to display any or all of the following information in the graph below: AP Load, 2.4GHz Load, 5GHz Load, and Channel. The graph section will display the load percentages for each of the selected variables over a period of time which can be altered using the slider below the upper graph.
Client Request Events	The Client Request Events displays the Time, Client name, Capability, State, WLAN and Requested Channels for all client request events on the Access Point.

Wireless Controller Statistics

The **Wireless Controller** screen displays information about peer controllers. As members of a cluster, a controller manages its own network and is ready to assume the load of an offline peer.

The **Wireless Controller** screen displays detailed statistics which include controller health, inventory of devices, wireless clients, adopted APs, rogue APs and WLANs. For more information, refer to the following:

- [Device Health](#)
- [Inventory](#)
- [Device](#)
- [Cluster Peers](#)
- [Adopted AP Statistics](#)
- [AP Detection](#)
- [Wireless Clients](#)
- [Wireless LANs](#)
- [Critical Resource](#)
- [Radios](#)
- [Interfaces](#)
- [Network](#)
- [DHCP Server](#)
- [Firewall](#)
- [IPsec](#)
- [Viewing Certificate Statistics](#)
- [Controller WIPS Statistics](#)
- [Advanced WIPS](#)
- [Sensor Server](#)
- [Captive Portal Statistics](#)

Device Health

[Wireless Controller Statistics](#)

The Device Health screen displays details such as hostname, device name, RF Domain name, radio RF quality and client RF quality.

To view the controller device health:

1. Select the **Statistics** tab from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Health** from the left-hand side of the UI.

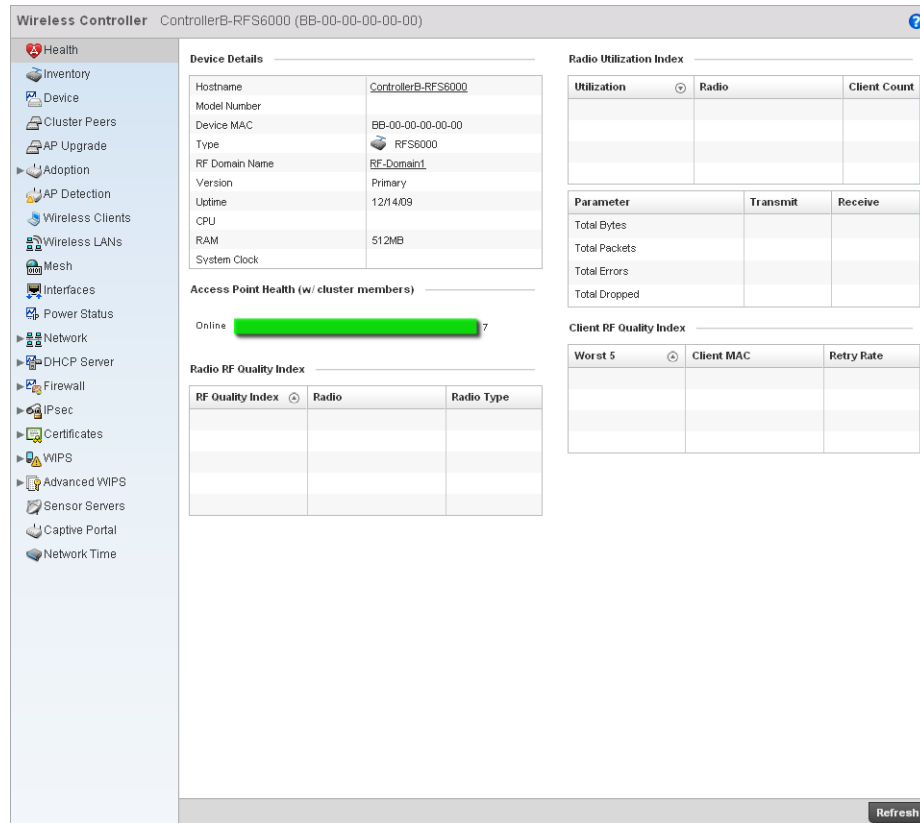


FIGURE 435 Wireless Controller Health screen

4. The **Device Details** field displays the following:

- Hostname** Displays the hostname of the wireless controller.
- Device MAC** Displays the MAC address of the controller.
- Type** Displays the controller type (RFS7000, RFS6000 or RFS4000).
- RF Domain Name** Displays the controller’s domain name.
- Version** Displays the version of the image running on the controller.
- Uptime** Displays the cumulative time since the controller was last rebooted or lost power.
- CPU** Displays the processor name.
- RAM** Displays the CPU memory in use.
- System Clock** Displays the system clock information.

5. The **Access Point Health (w/ cluster members)** field displays a bar chart showing how many Access Points are online and how many are offline. These are APs directly managed by the wireless controller. This data does not include Access Points associated to other controllers in the same cluster.
6. The **Radio RF Quality Index** field displays RF quality (overall effectiveness of the RF environment). Use this table to assess radio performance for improvement ideas.

The **RF Quality Index** field displays the following:

RF Quality Index	Displays the five radios with the lowest average quality.
Radio	Displays the hardware encoded MAC address of the radio.
Radio Type	Displays the radio type used by this Access Point.

7. The **Radio Utilization Index** field displays the following:

Utilization	Displays the traffic utilization indices of access points. The traffic utilization index measures how efficiently the traffic medium is used. It's defined as the percentage of the current throughput relative to the maximum relative possible throughput.
Radio	Displays the hardware encoded MAC address of the radio.
Client Count	Displays the number of clients associated with this Access Point.

Parameter	Displays the statistics in number of packets for: <ul style="list-style-type: none"> • Total Bytes - The total number of bytes that passed through the Access Point. • Total Packets - The total number of packets that passed through the Access Point. • Total Errors - The total error packets. • Total Dropped - The total dropped packets.
Transmit	Displays the total number of packets transmitted by the radio.
Receive	Displays the total number of packets received by the radio.

8. The **Client RF Quality Index** field displays the RF quality of the clients. Use this table to troubleshoot radios not optimally performing:

Worst 5	Displays the five radios having the lowest quality indices.
Client MAC	Displays the MAC address of the client.
Retry Rate	Displays the retry rate.

Inventory

[Wireless Controller Statistics](#)

The inventory statistics screen displays device types, radio types and wireless client on this controller. To view the controller inventory:

1. Select the **Statistics** tab from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Inventory** from the left-hand side of the UI.

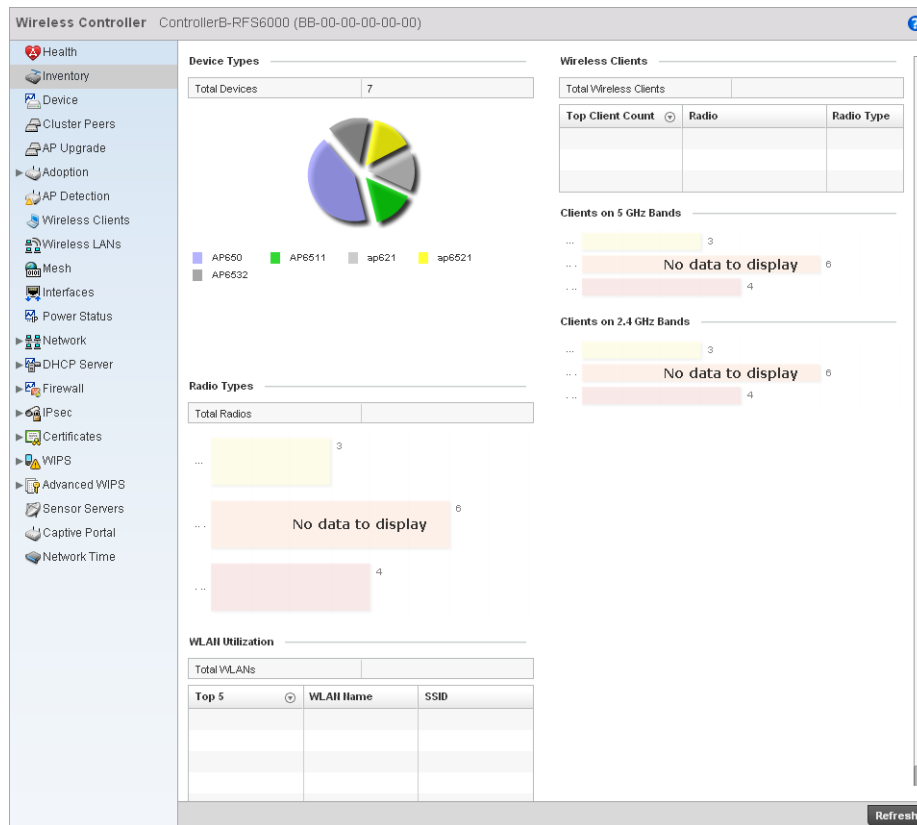


FIGURE 436 Wireless Controller Inventory screen

4. The **Device Types** table displays the total number of devices. An exploded pie chart depicts the distribution of the devices that are members of the managed network.
5. The **Radio Types** table displays the total number of radios used by the Access Points and wireless clients.
6. The **WLAN Utilization** table displays the total number of WLANs. This area also displays the following:

Top 5

Displays the traffic utilization index, which measures how efficiently the traffic medium is supported. It's defined as the percentage of current throughput relative to maximum possible throughput. Traffic indices include:

- 0–20 (very low utilization)
- 20–40 (low utilization)
- 40–60 (moderate utilization)
- 60 and above (high utilization)

WLAN Name

Displays the name of the WLAN.

SSID

Displays the Service Set ID associated with the WLAN.

7. The **Wireless Clients** table displays the total number of wireless clients associated with the controller. It displays the following:

Top Client Count	Displays clients with the highest detected traffic throughput.
Access Point	Displays the client's associated Access Point.
Radio Type	Displays the radio type associated with the Access Point.

8. The **Clients on 5 GHz Channels** field displays clients using radios in the 5 GHz frequency band.
9. The **Clients on 2.4 GHz Channels** field displays clients using radios in the 2.4 GHz frequency band.

Device

Wireless Controller Statistics

The Device Statistics screen provides detailed information about the selected device.

To view controller device statistics:

1. Select the **Statistics** tab from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Device** from the left-hand side of the UI.

The screenshot shows the 'Wireless Controller' interface for ControllerB-RFS6000 (BB-00-00-00-00-00). The left navigation pane includes sections like Health, Inventory, Device, Cluster Peers, AP Upgrade, Adoption, AP Detection, Wireless Clients, Wireless LANs, Mesh, Interfaces, Power Status, Network, DHCP Server, Firewall, IPsec, Certificates, WIPS, Advanced WIPS, Sensor Servers, Captive Portal, and Network Time. The main content area is divided into several sections:

- System:** A table with fields: Model Number, Version (Primary), Boot Partition (Primary), Fallback Enabled (Disabled), Fallback Image Triggered (False), and Next Boot (Primary).
- Licenses:** A table with fields: AAP Licenses, AAP Adoptions, AAP License, AP Licenses, AP Adoptions, and AP License.
- Additional Licenses:** A table with fields: ADSEC and WIPS.
- Firmware Images:** A table with fields: Primary Build Date (12/14/09), Primary Install Date (12/14/09), Primary Version (Primary), Secondary Build Date (12/14/09), Secondary Install Date (12/14/09), and Secondary Version (Secondary).
- Upgrade Status:** A table with fields: Upgrade Status (Successful) and Upgrade Status Time (Jan 13).
- DNS Mappings:** A table with columns: Name Server and Type.

A 'Refresh' button is located at the bottom right of the screen.

FIGURE 437 Wireless Controller Device screen

4. The **System** field displays the following:

Model Number	Displays the model number for the selected controller.
Version	Displays the unique alphanumeric firmware version name for the controller firmware.
Boot Partition	Displays the boot partitioning type.
Fallback Enabled	Displays whether fallback is enabled. The fallback feature enables a user to store both a legacy and new firmware version in memory. You can test the new software and use an automatic fallback mechanism, which loads the old version, if the new version fails.
Fallback Image Triggered	Displays whether the fallback image has been triggered. The fallback is a legacy software image stored in device memory. This allows a user to test a new version and revert to the older version if needed.
Next Boot	Designates this version as the version used the next time the controller is booted.

5. The **Firmware Images** field displays the following:

Primary Build Rate	Displays the build date when this version was created.
Primary Install Date	Displays the date this version was installed on the controller.
Primary Version	Displays the primary version string.
Secondary Build Date	Displays the build date when this secondary version was created.
Secondary Install Date	Displays the date this secondary version was installed on the controller.
Secondary Version	Displays the secondary version string.

6. The **Upgrade Status** field displays firmware upgrade statistics. The table provides the following:

Upgrade Status	Displays whether the image upgrade was successful.
Upgrade Status Time	Displays the time of the upgrade.

7. The **Licenses** field displays the following:

AAP Licenses	Displays the number of adaptive AP licenses on the controller. The maximum number permitted varies by controller platform.
AAP Adoptions	Displays the number of adaptive APs adopted by this controller.
AAP License	Displays the license string of the adaptive AP.
AP Licenses	Displays the number of AP licenses currently available on the controller. This value represents the maximum number of licenses the controller can adopt.
AP Adoptions	Displays the number of Access Points adopted by this controller.
AP License	Displays the license string of the AP.

8. The **Additional Licenses** area displays the following information:

ADSEC	Displays the number of Advanced Security licenses. This enables the Role Based firewall and increases the number of IP Sec VPN tunnels. The maximum number of IP Sec VPN tunnels varies by controller platform.
WIPS	Displays the number of WIPS licenses.

9. The **DNS Mappings** table displays the following:

Name Server	Displays any custom Name Server mappings on the controller.
Type	Displays the type of DNS mapping, if any, on the controller.

Cluster Peers

Wireless Controller Statistics

The cluster peer statistics screen provides cluster member information.

To view the controller cluster peer statistics:

1. Select the **Statistics** tab from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Cluster Peers** from the left-hand side of the UI.

Wireless Controller	MAC Address	Type	RF Domain Name	Online	Version
Controller A-RFS7000	AA-00-00-00-00-00	RFS7000	RF-Domain1	✓	primary
Controller C-RFS4000	CC-00-00-00-00-00	RFS4000	RF-Domain3	✓	primary

FIGURE 438 Wireless Controller Cluster Peers screen

4. The **Cluster Peers** screen displays the following:

Wireless Controller	Displays the IP addresses of current cluster members.
MAC Address	Displays the MAC address cluster members.
Type	Displays the type of cluster peer controller (RFS6000, RFS4000 etc.).

RF Domain Name	Displays each member's RF Domain.
Online	Displays whether a controller is online. If a controller is online a green check mark will be displayed, if it is offline a red X will display.
Version	Displays the each controller member's firmware version.

Adopted AP Statistics

Wireless Controller Statistics

The adopted AP statistics screen displays details about adopted APs.

To view adopted AP statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Adoption > Adopted APs** from the left-hand side of the UI.

Access Point	Type	RF Domain Name	Serial Number	Version	Adopted by	Adoption Time
BB-11-00-00-00-00	AP650	RF-Domain3	bb1100000000	version1		
BB-22-00-00-00-00	AP650	RF-Domain3	bb2200000000	version1		
BB-33-00-00-00-00	AP650	RF-Domain3	bb3300000000	version1		

FIGURE 439 Wireless Controller Adopted APs screen

4. The Adopted APs screen displays the following:

Access Point	Displays the name assigned to the Access Point.
AP MAC Address	Displays the hardcoded MAC address assigned to the unit when manufactured.
Type	Lists the AP model type.
RF Domain Name	Displays the Access Point's RF Domain assignment.
Online	Displays whether the listed AP is currently online and in service within the managed network.
Serial Number	Displays the Access Point's serial number. This is used for controller management.
Version	Displays the software (firmware) version used by the Access Point.

AP Adoption History

[Wireless Controller Statistics](#)

To view adopted AP Adoption History statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **AP Adoption History** from the left-hand side of the UI.

Event Name	AP MAC Address	Reason	Event Time
adopted	00-23-68-85-0F-84	N.A.	Wed May 11 2011 08:58:57 PM
adopted	00-23-68-86-45-38	N.A.	Wed May 11 2011 08:59:01 PM
adopted	00-23-68-9E-51-70	N.A.	Wed May 11 2011 08:59:45 PM
adopted	00-23-68-86-45-08	N.A.	Wed May 11 2011 09:02:19 PM
adopted	00-23-68-31-1A-94	N.A.	Wed May 11 2011 08:58:53 PM
adopted	00-23-68-86-45-8C	N.A.	Wed May 11 2011 09:02:26 PM
adopted	00-23-68-86-45-8C	N.A.	Wed May 11 2011 08:52:54 PM
adopted	00-23-68-41-DC-50	N.A.	Wed May 11 2011 08:52:51 PM
adopted	00-23-68-85-90-24	N.A.	Wed May 11 2011 08:52:42 PM
adopted	00-23-68-9E-51-70	N.A.	Wed May 11 2011 08:53:13 PM
adopted	00-23-68-86-47-BC	N.A.	Wed May 11 2011 09:03:35 PM
adopted	00-23-68-3A-71-F8	N.A.	Wed May 11 2011 08:53:19 PM
adopted	00-23-68-0F-C6-F8	N.A.	Wed May 11 2011 09:01:27 PM
adopted	00-23-68-86-45-44	N.A.	Wed May 11 2011 08:59:28 PM
adopted	00-23-68-86-45-38	N.A.	Wed May 11 2011 08:52:44 PM
adopted	00-23-68-85-8F-84	N.A.	Wed May 11 2011 08:52:44 PM
adopted	00-23-68-31-1A-08	N.A.	Wed May 11 2011 08:52:53 PM
adopted	00-23-68-86-44-B4	N.A.	Wed May 11 2011 09:01:58 PM
adopted	00-23-68-86-47-C0	N.A.	Wed May 11 2011 08:59:23 PM
adopted	00-23-68-31-1A-08	N.A.	Wed May 11 2011 09:01:54 PM
adopted	00-23-68-86-44-B8	N.A.	Wed May 11 2011 08:59:03 PM
adopted	00-23-68-85-92-80	N.A.	Wed May 11 2011 08:52:45 PM
adopted	00-23-68-86-45-5C	N.A.	Wed May 11 2011 08:52:45 PM
adopted	00-23-68-86-44-B8	N.A.	Wed May 11 2011 08:52:44 PM
adopted	00-23-68-86-47-C0	N.A.	Wed May 11 2011 08:52:45 PM
adopted	00-23-68-0F-44-24	N.A.	Wed May 11 2011 08:56:52 PM
adopted	00-23-68-31-1A-94	N.A.	Wed May 11 2011 08:52:44 PM
adopted	00-23-68-85-92-80	N.A.	Wed May 11 2011 08:58:58 PM
adopted	00-23-68-86-45-08	N.A.	Wed May 11 2011 08:52:49 PM
adopted	5C-0E-8B-08-44-F8	N.A.	Wed May 11 2011 08:53:03 PM
adopted	00-23-68-3A-71-F8	N.A.	Wed May 11 2011 09:00:03 PM
adopted	00-23-68-86-45-48	N.A.	Wed May 11 2011 09:02:31 PM
adopted	00-23-68-0F-44-24	N.A.	Wed May 11 2011 08:52:46 PM
adopted	00-23-68-86-45-5C	N.A.	Wed May 11 2011 08:58:58 PM
adopted	5C-0E-8B-08-44-F8	N.A.	Wed May 11 2011 08:04:49 PM

FIGURE 440 AP Adoption History screen

4. The **Adopted Devices** screen provides the following

- Event Name** Displays the adoption status of each AP as either adopted or un-adopted.
- AP MAC Address** Displays the Media Access Control (MAC) address of each Access Point that the controller has attempted to adopt.
- Reason** Displays the reason code for each event listed in the adoption history statistics table.
- Event Time** Displays day, date and time for each Access Point adoption attempt.

Pending Adoptions

Wireless Controller Statistics

The *Adopted Devices* screen displays a list of devices adopted to the wireless controller managed network. Use this screen to view a list of devices and their current status.

To view adopted AP statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Pending Adoptions** from the left-hand side of the UI.

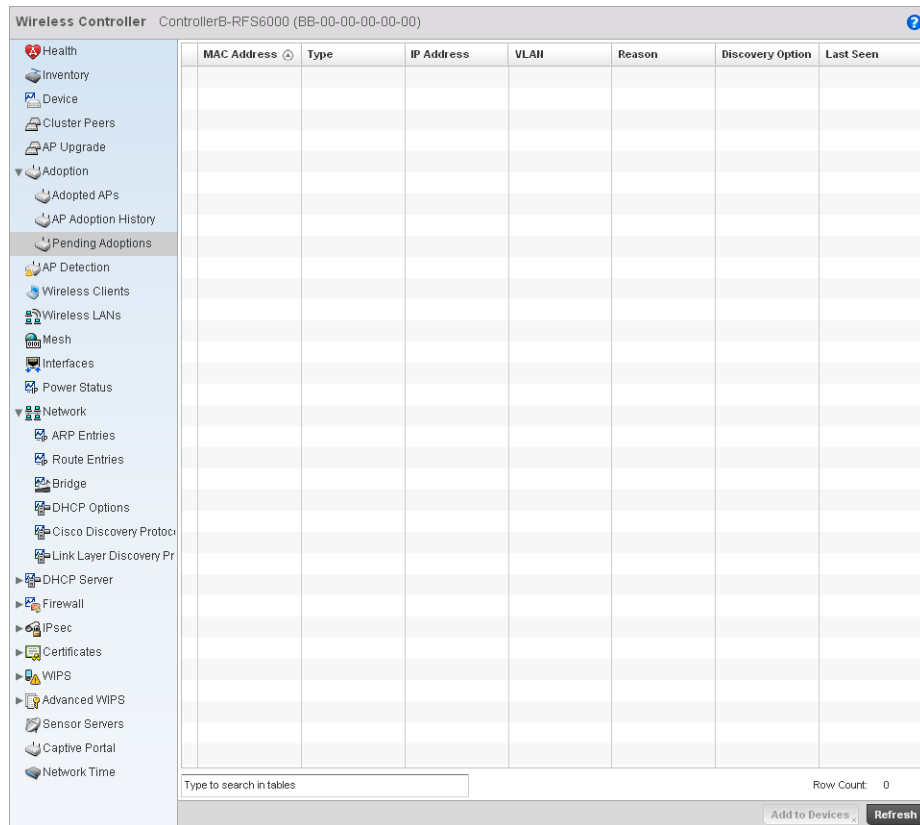


FIGURE 441 Pending Adoptions Devices screen

4. The **Adopted Devices** screen provides the following

MAC Address	Displays the MAC address of the device pending adoption.
Type	Displays the AP type (either AP650, AP7131, or AP6511).
IP Address	Displays the current IP Address of the device pending adoption.
VLAN	Displays the current VLAN number of the device pending adoption.
Reason	Displays the status as to why the device is still pending adoption.
Discovery Option	Displays the discovery option code for each AP listed pending adoption.
Last Seen	Displays the date and time stamp of the last time the device was seen. Click the arrow next to the date and time to toggle between standard time and UTC.

AP Detection

[Wireless Controller Statistics](#)

The AP detection statistics screen displays details about detected rogue Access Points.

To view the controller AP detection statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **AP Detection** from the left-hand side of the UI.

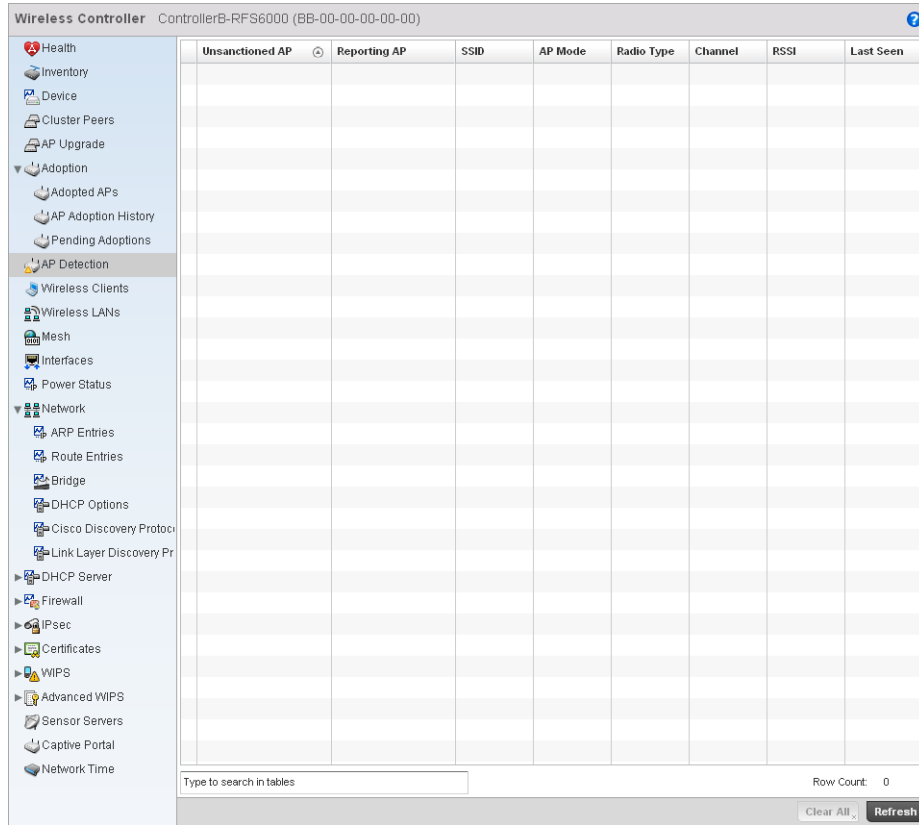


FIGURE 442 Wireless Controller AP Detection screen

4. The AP Detection screen displays the following:

- Unsanctioned AP** Displays the MAC address of unsanctioned APs.
- Reporting AP** Displays a numerical value for the radio used with the detecting AP.
- SSID** Displays the SSID of each unsanctioned AP.
- AP Mode** Displays the mode of the unsanctioned device (either Access Point or wireless client).
- Radio Type** Displays the unsanctioned AP's radio type. The radio can be 802.11b, 802.11bg, 802.11g, 802.11a or 802.11an.
- Channel** Displays the unsanctioned AP's current operating channel.
- RSSI** Displays the Received Signal Strength Indicator (RSSI) for rogue APs.
- Last Seen** Displays when the unsanctioned AP was last seen by the detecting AP.

Wireless Clients

[Wireless Controller Statistics](#)

The wireless client statistics screen displays details about wireless clients.

To view the wireless client statistics of the controller:

1. Select the **Statistics** menu from the Web UI.

2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Wireless Clients** from the left-hand side of the UI.

Client MAC	WLAN	Username	State	VLAN	IP Address	Vendor
AA-11-11-00-00-00	wlan1	user1	Roaming	1	10.1.1.1	Motorola
AA-11-22-00-00-00	wlan1	user1	associating	2,100	10.1.1.1	Motorola

FIGURE 443 Wireless Controller Wireless Clients screen

4. The **Wireless Clients** screen displays the following:

Client MAC	Displays the MAC address of the wireless client.
WLAN	Displays the name of the WLAN the client is currently associated with. Use this information to determine if the client/WLAN placement best suits the intended operation and the client's coverage area.
Username	Displays the unique name of the administrator or operator.
State	Displays whether the client is online or offline.
VLAN	Displays the name of the client's current VLAN mapping.
IP Address	Displays the unique IP address of the client. Use this address as necessary throughout the applet for filtering and device intrusion recognition and approval.
Vendor	Displays the name of the client vendor.

Wireless LANs

[Wireless Controller Statistics](#)

The wireless LAN statistics screen displays performance statistics for each WLAN. Use this information to assess if configuration changes are required to improve network performance.

To view the wireless LAN statistics of the controller:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Wireless LANs** from the left-hand side of the UI.

The screenshot shows the 'Wireless LANs' screen for ControllerA-RFS7000 (AA-00-00-00-00). The left navigation pane is expanded to 'Wireless LANs'. The main table contains the following data:

WLAN Name	SSID	Traffic Index	Radio Count	Tx Bytes	Tx User Data Rate	Rx Bytes	Rx User Data Rate
WLAN3	WLAN3	21 (Low)	4	3,001	4,001 kbps	2,001	10,011 kbps
WLAN4	WLAN4	75 (High)	31	2,300	4,100 kbps	2,100	10 kbps

At the bottom of the table, there is a search bar with the text 'Type to search in tables' and a 'Refresh' button. The 'Row Count' is displayed as 2.

FIGURE 444 Wireless Controller Wireless LANs screen

4. The WLAN screen displays the following:

- WLAN Name** Displays the name of the WLAN the controller is currently utilizing.
- SSID** Displays the Service Set ID associated with each WLAN.
- Traffic Index** Displays the traffic utilization index, which measures how efficiently the traffic medium is used. It's defined as the percentage of current throughput relative to the maximum possible throughput. Traffic indices are:
 - 0–20 (very low utilization)
 - 20–40 (low utilization)
 - 40–60 (moderate utilization)
 - 60 and above (high utilization)
- Radio Count** Displays the number of radios associated with this WLAN.
- Tx Bytes** Displays the data transmitted in bytes on the selected WLAN.

- Tx User Data Rate** Displays the average user data rate.
- Rx Bytes** Displays the data received in bytes on the selected WLAN.
- Rx User Data Rate** Displays the average user data rate.

Critical Resource

Wireless Controller Statistics

The Critical Resources statistics screen displays a list of device IP addresses on the network (gateways, routers etc.). These defined IP address is critical to the health of the access point managed network. These device addresses are pinged regularly by the access point. If there is a connectivity issue, an event is generated stating a critical resource is unavailable.

To view the Critical Resource statistics of the controller:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Critical Resource** from the left-hand side of the UI.

IP Address	VLAN	Ping Mode	State

Type to search in tables Row Count: 0

Refresh

FIGURE 445 Wireless Controller Critical Resource screen

4. The **Critical Resource** screen displays the following:

IP Address	Lists the IP address of the critical resource. This is the address the device assigned and is used by the access point to ensure the critical resource is available.
VLAN	Lists the access point VLAN on which the critical resource is available.
Ping Mode	Describes the ping mode as either: arp-only - Uses ARP for only pinging the critical resource. ARP is used to resolve hardware addresses when only the network layer address is known. OR arp-icmp - Uses both ARP and ICMP for pinging the critical resource and sending control messages (device not reachable, requested service not available, etc.).
State	Defines the operational state of each listed critical resource.

Radios

[Wireless Controller Statistics](#)

The radio statistics screen provides radio association data, including radio ID, connected APs, radio type, quality index and *Signal to Noise Ratio* (SNR).

To view the radio statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Radio** from the left-hand side of the UI.

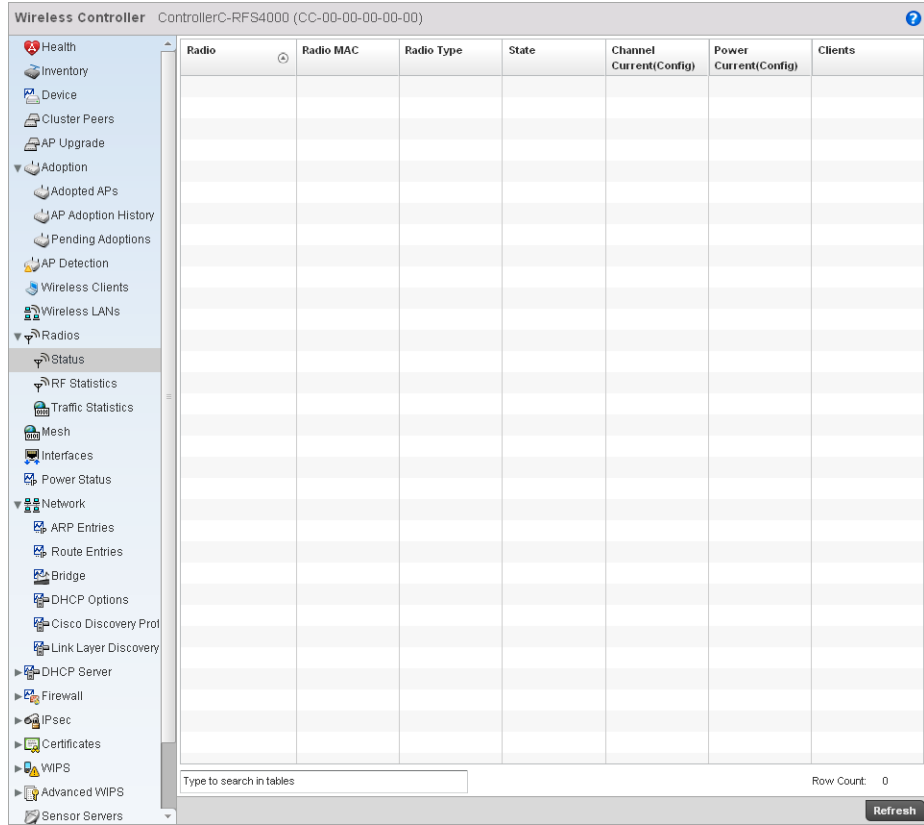
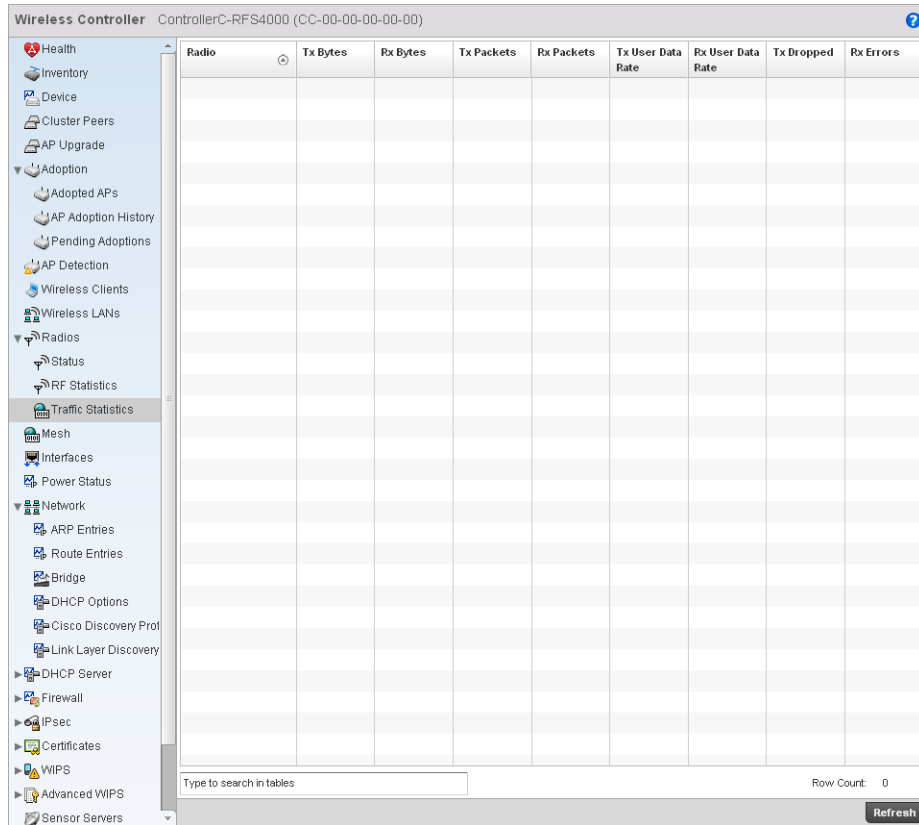


FIGURE 446 Wireless Controller Radio Status screen



4.

[Wireless Controller Statistics](#)

5. Select a **Wireless Controller** node from the left navigation pane.

The screenshot shows the 'Wireless Controller' interface for 'ControllerA-RFS7000 (AA-00-00-00-00-00)'. The left sidebar contains a navigation tree with categories like Health, Inventory, Device, Cluster Peers, AP Upgrade, Adoption, AP Detection, Wireless Clients, Wireless LANs, Mesh, Interfaces, Power Status, Network, ARP Entries, Route Entries, Bridge, DHCP Options, Cisco Discovery Prot, Link Layer Discovery, DHCP Server, Firewall, IPsec, Certificates, WIPS, Advanced WIPS, Sensor Servers, Captive Portal, and Network Time. The main content area is titled 'ge1' and has two tabs: 'General' and 'Network Graph'. The 'General' tab is active and displays the following data:

General	
Name	ge1
Interface MAC Address	11-11-00-11-00-11
IP Address	1.2.3.4
IP Address Type	
Secondary IPs	1 item
Hardware Type	ethernet
Index	1
Access Setting	Access
Access VLAN	
Native VLAN	
Tagged Native VLAN	
Allowed VLANs	
Administrative Status	true

Specification	
Media Type	one
Protocol	
MTU	1,000
Mode	Layer 2
Metric	
Maximum Speed	100M
Admin Speed	50
Operator Speed	true
Admin Duplex Setting	Half
Current Duplex Setting	Full

Traffic	
Good Octets Sent	2,900
Good Octets Received	14,100
Good Pkts Sent	2,009
Good Pkts Received	24,020

Errors	
Bad Pkts Received	47,040
Collisions	9,200
Late Collisions	3,500
Excessive Collisions	6,200
Drop Events	3,002
Tx Undersize Pkts	5,500
Over-size Pkts	56,400
MAC Transmit Error	2,100
MAC Receive Error	8,200
Bad CRC	2,700

Receive Errors	
Rx Frame Errors	11
Rx Length Errors	12
Rx FIFO Errors	10
Rx Missed Errors	13
Rx Over Errors	14

Transmit Errors	
Tx Errors	18
Tx Dropped	17
Tx Aborted Errors	15
Tx Carrier Errors	16
Tx FIFO Errors	19
Tx Heartbeat Errors	20
Tx Window Errors	21

FIGURE 447 Wireless Controller Interfaces screen

The *Interface Statistics* screen can be divided into:

- [Viewing Interface Details](#)
- [Viewing Interface Statistics Graph](#)

Viewing Interface Details

Interfaces

The **General** table displays the following:

Name	Displays the name of the interface.
Interface MAC Address	Displays the MAC address of the interface.
IP Address	IP address of the interface.
IP Address Type	Lists the interface's IP address.
Hardware Type	Displays the networking technology.
Index	Displays the unique numerical identifier for the interface.
Access VLAN	Displays the tag assigned to the native VLAN.
Access Setting	Displays the VLAN mode as either <i>Access</i> or <i>Trunk</i> .
Administrative Status	Displays whether the interface is currently UP or DOWN.

The **Specification** table displays the following information:

Media Type	Displays the physical connection type of the interface. Medium types include: <i>Copper</i> - Used on RJ-45 Ethernet ports <i>Optical</i> - Used on fibre optic gigabit Ethernet ports
Protocol	Displays the routing protocol used by the interface.
MTU	Displays the <i>maximum transmission unit</i> (MTU) setting configured on the interface. The MTU value represents the largest packet size that can be sent over a link. 10/100 Ethernet ports have a maximum setting of 1500.
Mode	The mode can be either: <i>Access</i> – The Ethernet interface accepts packets only from native VLANs. <i>Trunk</i> – The Ethernet interface allows packets from a list of VLANs you can add to the trunk.
Metric	Displays the metric associated with the interface's route.
Maximum Speed	Displays the maximum speed the interface uses to transmit or receive data.
Admin Speed	Displays the speed the port can transmit or receive. This value can be either <i>10</i> , <i>100</i> , <i>1000</i> or <i>Auto</i> . This value is the maximum port speed in Mbps. Auto indicates the speed is negotiated between connected devices.
Operator Speed	Displays the current speed of data transmitted and received over the interface.
Admin Duplex Setting	Displays the administrator's duplex setting.
Current Duplex Setting	Displays the interface as either <i>half duplex</i> , <i>full duplex</i> or <i>unknown</i> .

The **Traffic** table displays the following:

Good Octets Sent	Displays the number of octets (bytes) with no errors sent by the interface.
Good Octets Received	Displays the number of octets (bytes) with no errors received by the interface.
Good Packets Sent	Displays the number of good packets transmitted.
Good Packets Received	Displays the number of good packets received.
Mcast Pkts Sent	Displays the number of multicast packets sent through the interface.
Mcast Pkts Received	Displays the number of multicast packets received through the interface.
Bcast Pkts Sent	Displays the number of broadcast packets sent through the interface.
Bcast Pkts Received	Displays the number of broadcast packets received through the interface.
Packet Fragments	Displays the number of packet fragments transmitted or received through the interface.
Jabber Pkts	Displays the number of packets transmitted through the interface larger than the MTU.

The **Errors** table displays the following:

Bad Pkts Received	Displays the number of bad packets received through the interface.
Collisions	Displays the number of collisions.
Late Collisions	A late collision is any collision that occurs after the first 64 octets of data have been sent. Late collisions are not normal, and usually the result of out of specification cabling or a malfunctioning device.
Excessive Collisions	Displays the number of excessive collisions. Excessive collisions occur when the traffic load increases to the point a single Ethernet network cannot handle it efficiently.
Drop Events	Displays the number of dropped packets transmitted or received through the interface.
Tx Undersize Pkts	Displays the number of undersized packets transmitted through the interface.
Oversize Pkts	Displays the number of oversized packets transmitted through the interface.
MAC Transmit Error	Displays the number of failed transmits due to an internal MAC sublayer error that's not a late collision, due to excessive collisions or a carrier sense error.
MAC Receive Error	Displays the number of received packets that failed due to an internal MAC sublayer that's not a late collision, an excessive number of collisions or a carrier sense error.
Bad CRC	Displays the CRC error. The CRC is the 4 byte field at the end of every frame. The receiving station uses it to interpret if the frame is valid. If the CRC value computed by the interface does not match the value at the end of frame, it is considered as a bad CRC.

The **Receive** table displays the following:

Rx Frame Errors	Displays the number of frame errors received at the interface. A frame error occurs when data is received, but not in an expected format.
Rx Length Errors	Displays the number of length errors received at the interface. Length errors are generated when the received frame length was either less or over the Ethernet standard.
Rx FIFO Errors	Displays the number of FIFO errors received at the interface. First-in First-out queueing is an algorithm that involves buffering and forwarding of packets in the order of arrival. FIFO entails no priority. There is only one queue, and all packets are treated equally. An increase in FIFO errors indicates a probable hardware malfunction.
Rx Missed Errors	Displays the number of missed packets. Packets are missed when the hardware received FIFO has insufficient space to store an incoming packet.
Rx Over Errors	Displays the number of overflow errors received. Overflows occur when a packet size exceeds the allocated buffer size.

The **Transmit Errors** field displays the following:

Tx Errors	Displays the number of packets with errors transmitted on the interface.
Tx Dropped	Displays the number of transmitted packets dropped from the interface.
Tx Aborted Errors	Displays the number of packets aborted on the interface because a clear-to-send request was not detected.
Tx Carrier Errors	Displays the number of carrier errors on the interface. This generally indicates bad Ethernet hardware or bad cabling.

Tx FIFO Errors	Displays the number of FIFO errors transmitted at the interface. <i>First-in First-Out</i> queueing is an algorithm that involves the buffering and forwarding of packets in the order of arrival. FIFO uses no priority. There is only one queue, and all packets are treated equally. An increase in the number of FIFO errors indicates a probable hardware malfunction.
Tx Heartbeat Errors	Displays the number of heartbeat errors. This generally indicates a software crash, or packets stuck in an endless loop.
Tx Window Errors	Displays the number of window errors transmitted. TCP uses a sliding window flow control protocol. In each TCP segment, the receiver specifies the amount of additional received data (in bytes) the receiver is willing to buffer for the connection. The sending host can send only up to that amount. If the sending host transmits more data before receiving an acknowledgment, it constitutes a window error.

Viewing Interface Statistics Graph

Interfaces

The *Network Graph* tab displays interface statistics the controller continuously collects for interface statistics. Even when the interface statistics graph is closed, data is still tallied. Display the interface statistics graph periodically for assessing the latest interface information.

To view a detailed graph for an interface, select an interface and drop it on to the graph. The graph displays Port Statistics as the Y-axis and the Polling Interval as the X-axis.

To view the Interface Statistics graph:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Interfaces** from the left-hand side of the UI, then select the **Network Graph** tab.

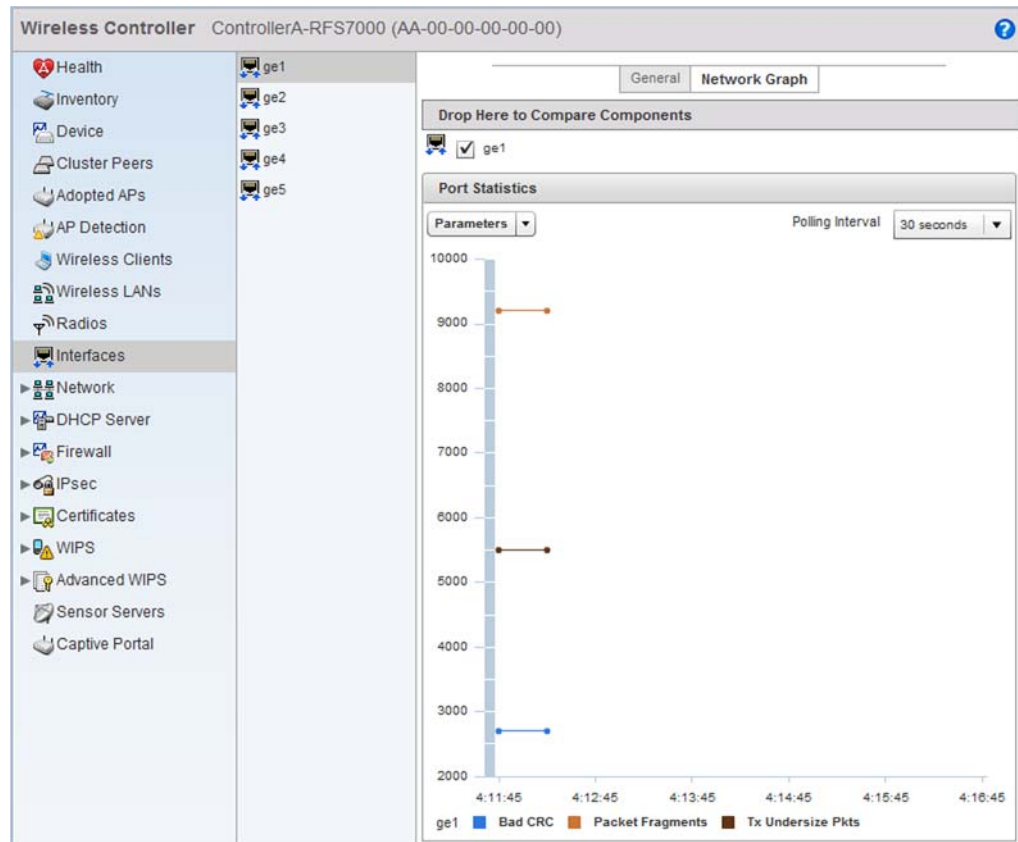
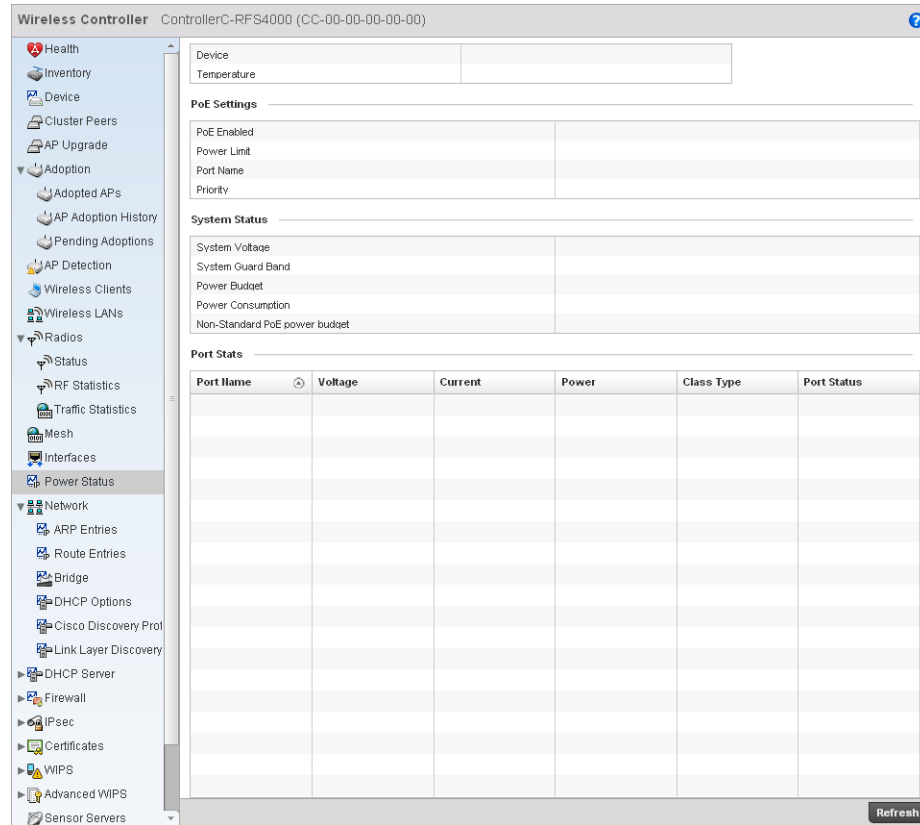


FIGURE 448 Wireless Controller Interfaces Network Graph screen

Wireless Controller Statistics

4. Select a **Wireless Controller** node from the left navigation pane.



Network

Wireless Controller Statistics

Use the *Network* screen to view information for ARP, DHCP, Routing and Bridging. Each of these screens provides enough data to troubleshoot issues related to the following:

- [ARP Entries](#)
- [Route Entries](#)
- [Bridge](#)
- [DHCP Options](#)
- [Cisco Discovery Protocol](#)
- [Link Layer Discovery Protocol](#)

ARP Entries

Network

The *Address Resolution Protocol* (ARP) is a networking protocol for determining a network host's hardware address when its IP address or network layer address is known.

To view the ARP entries on the network statistics screen:

1. Select the **Statistics** menu from the Web UI.

2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Network > ARP Entries** from the left-hand side of the UI.

IP Address	ARP MAC Address	Type	VLAN
10.0.0.5	55:11:22:33:44:55	ts	5

FIGURE 449 Wireless Controller Network ARP Entries screen

4. The ARP Entries screen displays the following:

IP Address	Displays the IP address of the client being resolved.
ARP MAC Address	Displays the MAC address of the device where an IP address is being resolved.
Type	Defines whether the entry was added statically or created dynamically in respect to network traffic. Entries are typically static.
VLAN	Displays the name of the VLAN where the IP address was found.

Route Entries

Network

The route entries screen displays data for routing packets to a defined destination. When an existing destination subnet does not meet the needs of the network, add a new destination subnet, subnet mask and gateway.

To view the route entries:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.

3. Select **Network > Route Entries** from the left-hand side of the UI.

Destination	FLAGS	Gateway	Interface
destination5	false	gw5	ge5

FIGURE 450 Wireless Controller Network Route Entries

4. The **Route Entries** screen provides the following information:

Destination	Displays the IP address of a specific destination address.
DKEY	<TBD>
FLAGS	Displays the flags for this route entry. <i>C</i> indicates a connected state. <i>G</i> indicates a gateway.
Gateway	Displays the gateway IP address used to route packets to the destination subnet.
Interface	Displays the name of the interface of the destination subnet.

Bridge

Network

Bridging is a forwarding technique making no assumption about where a particular network address is located. It depends on flooding and the examination of source addresses in received packet headers to locate unknown devices. Once a device is located, its location is stored in a table to avoid broadcasting to that device again. Bridging is limited by its dependency on flooding, and is used in local area networks only. A bridge and a controller are very similar, since a controller is a bridge with a number of ports.

The *Bridge* screen provides details about the *Integrated Gateway Server (IGS)*, which is a router connected to an access point. The IGS performs the following:

- *Issues IP addresses*
- *Throttles bandwidth*
- *Permits access to other networks*
- *Times out old logins*

This screen also provides information about the *Multicast Router (MRouter)*, which is a router program that distinguishes between multicast and unicast packets and how they should be distributed along the Multicast Internet. Using an appropriate algorithm, a multicast router tells a switching device what to do with the multicast packet.

This screen is partitioned into two tabs:

- [Details](#)
- [MAC Address](#)

Details

To view network bridge details:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Network > Bridge** from the left-hand side of the UI, and select the **Details** tab.

Wireless Controller ControllerC-RFS4000 (CC-00-00-00-00-00)

Details MAC Address

Integrated Gateway Server (IGS)

VLAN 1

Integrated Gateway Server (IGS)

Group Address	Port Members	Version
10.1.1.1	false	6,000

Multicast Router (MRouter)

Learn Mode	Port Members	Query Interval	Version	VLAN
✗ Disabled	1	1,800	6,000	1

Refresh

FIGURE 451 Wireless Controller Network Bridge Details screen

4. The **Integrated Gateway Server (IGS)** field displays the following:

VLAN	Displays the integrated gateway server's VLAN.
Group Address	Displays the IP Address for group broadcast on the bridge.
Port Members	Displays the ports included on this server. For example, <i>10/100 Base Tx Ethernet port, RS485 serial port</i> etc.
Version	Displays the gateway server version.

MAC Address

To view network bridge MAC address information:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Network > Bridge** from the from the left hand side of the UI, and select the **MAC Address** tab.

Bridge Name	MAC Address	Interface	VLAN	Forwarding
bridge1	11-22-33-44-55-66	eth0	1	true

FIGURE 452 Wireless Controller Network Bridge MAC Address screen

4. The **MAC Address** tab displays the following:

Bridge Name	Displays the name of the network bridge.
MAC Address	Displays the MAC address of each listed bridge.
Interface	Displays the interface the bridge uses to transfer packets.
VLAN	Displays the VLAN the bridge belongs to.
Forwarding	Displays whether the bridge is forwarding packets. A bridge can only forward packets, thus the display is either <i>true</i> or <i>false</i> .

DHCP Options

Network

The controller contains an internal *Dynamic Host Configuration Protocol* (DHCP) server. The DHCP server can provide the dynamic assignment of IP addresses automatically. This is a protocol that includes IP address allocation and delivery of host-specific configuration parameters from a DHCP server to a host. Some of these parameters include IP address, gateway and network mask.

To view network DHCP options:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.

3. Select **Network > DHCP Options**.

Server Information	Image File	Configuration	Legacy Adoption	Adoption
server_information_1	image_1	configuration_1		
server_information_2	image_2	configuration_2		

FIGURE 453 Wireless Controller Network DHCP Options screen

4. The **DHCP Options** screen describes the following:

Server Information Displays the name of the DHCP server.

Image File Displays the image file name. BOOTP or the bootstrap protocol can be used to boot diskless clients. An image file is sent from the boot server. The file contains the operating system image. DHCP servers can be configured to support BOOTP.

Configuration Displays the Configuration name for each DHCP Server.

Cluster Configuration Displays the Cluster Configuration name for each DHCP Server.

Cisco Discovery Protocol

Network

To view a network's CDP Statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Network > Cisco Discovery Protocol**.

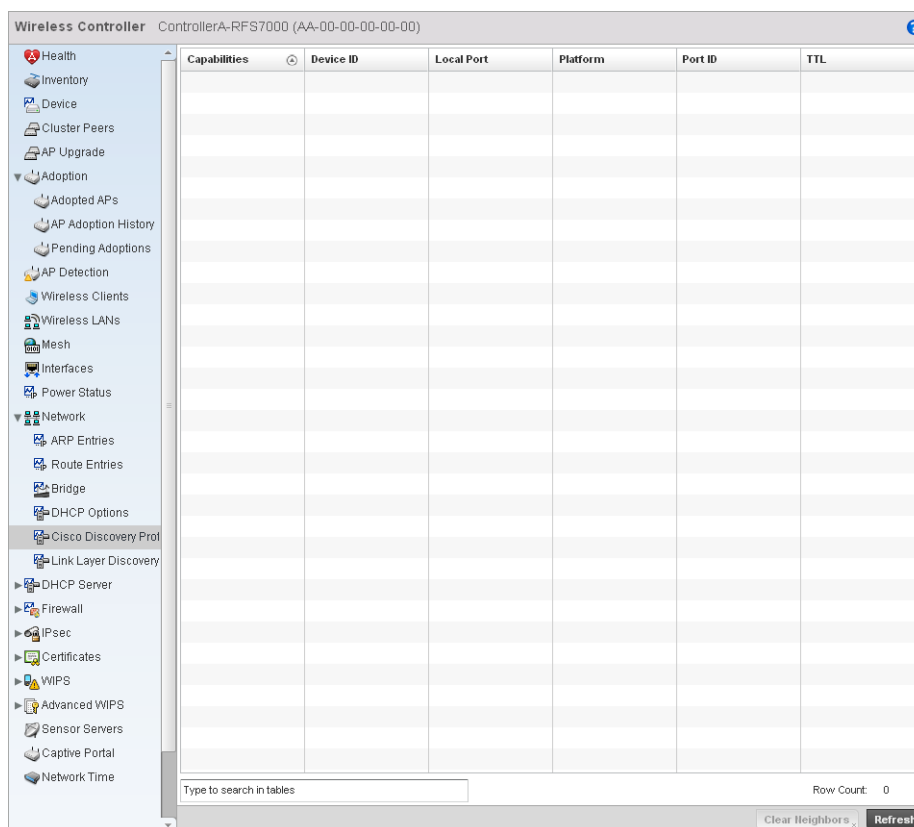


FIGURE 454 Access Point Network Cisco Discovery Protocol screen

4. The **Cisco Discovery Protocol** screen displays the following:

Capabilities	Displays the capabilities code for the device either Router, Trans Bridge, Source Route Bridge, Switch, Host, IGMP or Repeater.
Device ID	Displays the configured device ID or name for each device in the table.
Local Port	Displays the local port name for each CDP capable device.
Platform	Displays the model number of the CDP capable device.
Port ID	Displays the identifier for the local port.
TTL	Displays the time to live for each CDP connection.
Clear Neighbors	Click Clear Neighbors to remove all known CDP neighbors from the table.

Link Layer Discovery Protocol

Network

To view a network's Link Layer Discovery Protocol statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Network > Link Layer Discovery Protocol**.

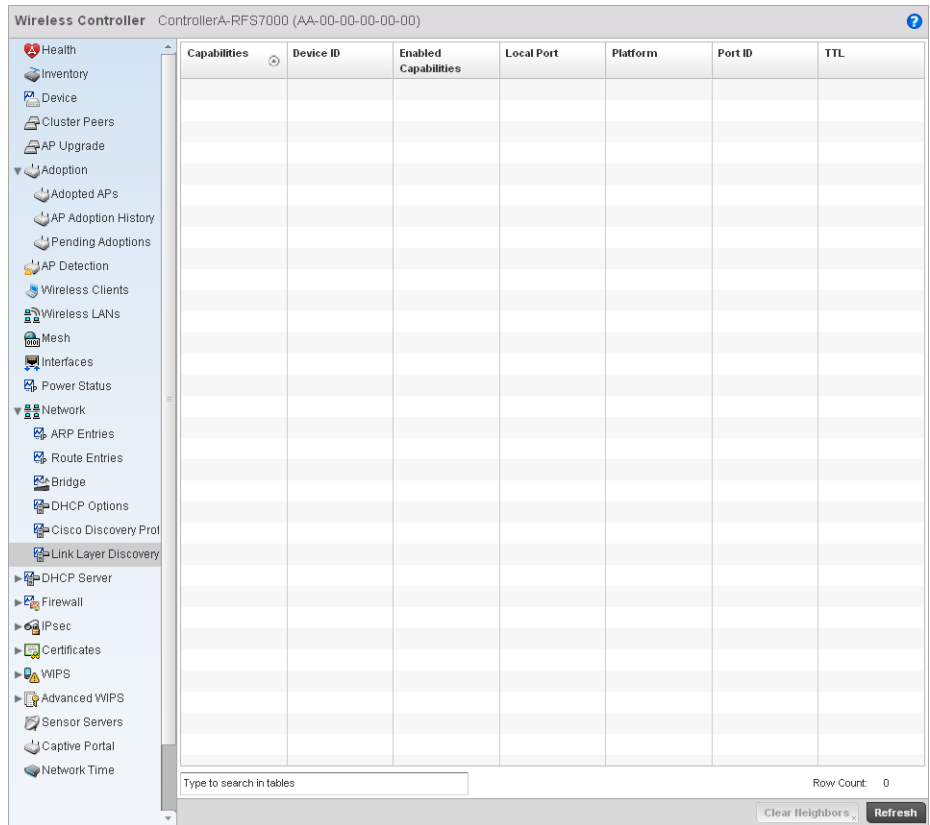


FIGURE 455 Access Point Link Layer Discovery screen

4. The Link Layer Discovery Protocol screen displays the following:

- Capabilities** Displays the capabilities code for the device either Router, Trans Bridge, Source Route Bridge, Switch, Host, IGMP or Repeater.
- Device ID** Displays the configured device ID or name for each device in the table.
- Enabled Capabilities** Displays which of the device capabilities are currently enabled.
- Local Port** Displays the local port name for each LLDP capable device.
- Platform** Displays the model number of the LLDP capable device.
- Port ID** Displays the identifier for the local port.
- TTL** Displays the time to live for each LLDP connection.
- Clear Neighbors** Click Clear Neighbors to remove all known LLDP neighbors from the table.

DHCP Server

Wireless Controller Statistics

The controller contains an internal *Dynamic Host Configuration Protocol* (DHCP) server. DHCP can provide IP addresses automatically. DHCP is a protocol that includes mechanisms for IP address allocation and delivery of host-specific configuration parameters (IP address, network mask gateway etc.) from a DHCP server to a host.

To review DHCP server statistics, refer to the following:

- [Viewing General DHCP Information](#)
- [Viewing DHCP Binding Information](#)
- [Viewing DHCP Server Networks Information](#)

Viewing General DHCP Information

DHCP Server

To view general DHCP information:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **DHCP Server > General** from the left-hand side of the UI.

Wireless Controller ControllerA-RFS7000 (AA-00-00-00-00)

Health
Inventory
Device
Cluster Peers
AP Upgrade
Adoption
Adopted APs
AP Adoption History
Pending Adoptions
AP Detection
Wireless Clients
Wireless LANs
Mesh
Interfaces
Power Status
Network
ARP Entries
Route Entries
Bridge
DHCP Options
Cisco Discovery Prot
Link Layer Discovery
DHCP Server
General
Bindings
Networks
Firewall
IPsec
Certificates
WIPS
Advanced WIPS
Sensor Servers
Captive Portal

Status

Interfaces	ge1
State	s1

DHCP Bindings

IP Address	Name
210.200.20.221	DDNS5

DHCP Manual Bindings

IP Address	Client Id
1.2.3.4	clientname
1.2.3.5	clientname2

Refresh

FIGURE 456 Wireless Controller DHCP Server screen

4. The **Status** table defines the following:

Interfaces	Displays the controller interface used for the created DHCP configuration.
State	Displays the current state of the DHCP server.

5. The **DDNS Bindings** table displays the following:

IP Address	Displays the IP address assigned to the client.
Name	Displays the domain name mapping corresponding to the listed IP address.

6. The **DHCP Manual Bindings** table displays the following:

IP Address	Displays the IP address for each client with a listed MAC address.
Client Id	Displays the MAC address (client hardware ID) of the client.

Viewing DHCP Binding Information

DHCP Server

The DHCP binding information screen displays DHCP binding information such as expiry time, client IP addresses and their MAC address.

To view the DHCP binding information:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **DHCP Server > Bindings** from the left-hand side of the UI.

Expiry Time	IP Address	DHCP MAC Address
20:17:45	160.132.30.31	10:00:00:ff:cc

FIGURE 457 Wireless Controller DHCP Binding screen

4. The **Bindings** screen displays the following:

Expiry Time	Displays the expiration of the lease used by the client for controller DHCP resources.
IP Address	Displays the IP address for clients whose MAC address is listed in the Client Id column.
DHCP MAC Address	Displays the client MAC address (ID) of the client.

Viewing DHCP Server Networks Information

DHCP Server

The DHCP server maintains a pool of IP addresses and client configuration parameters (default gateway, domain name, name servers etc). On receiving a valid client request, the server assigns the computer an IP address, a lease (the validity of time), and other IP configuration parameters.

The *Networks* screen provides network pool information such as the subnet for the addresses you want to use from the pool, the pool name, the used addresses and the total number of addresses.

To view the **DHCP Server Networks** information:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **DHCP Server > Networks** from the left-hand side of the UI.

Name	Subnet Address	Used Addresses	Total Addresses
gd1	13.31.13.13/24	12	11

FIGURE 458 Wireless Controller DHCP Network screen

4. The **Networks** screen displays the following:

Name	Displays the name of the network pool from which IP addresses can be issued to DHCP client requests on the current interface.
Subnet Address	Displays the subnet for the IP addresses used from the network pool.
Used Addresses	Displays the host IP addresses allocated by a DHCP server.
Total Addresses	Displays the total number of IP addresses in the network pool.

Firewall

[Wireless Controller Statistics](#)

A firewall is designed to block unauthorized access while permitting authorized communications. It's a device or a set of devices configured to permit or deny computer applications based on a set of rules. For more information, refer to the following:

- [Viewing Packet Flow Statistics](#)
- [Viewing Denial of Service Statistics](#)
- [IP Firewall Rules](#)
- [MAC Firewall Rules](#)
- [Viewing DHCP Snooping Statistics](#)

Viewing Packet Flow Statistics

Firewall

The *Packet Flows* screen displays data traffic packet flow utilization. The chart represents the different protocol flows supported, and displays a proportional view of the flows in respect to their percentage of data traffic utilized. The *Total Active Flows* field displays the total number of flows supported by the controller.

To view the packet flow statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Firewall > Packet Flows** from the left-hand side of the controller UI.

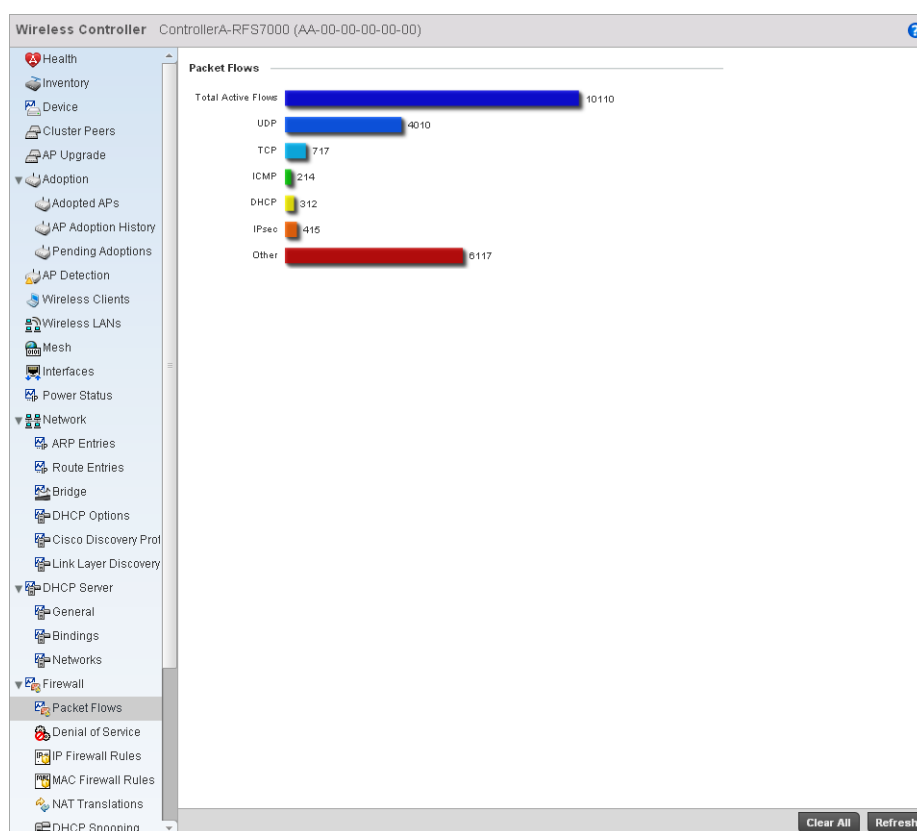


FIGURE 459 Wireless Controller Firewall Packet Flows screen

Viewing Denial of Service Statistics

Firewall

A *denial-of-service attack* (DoS attack), or distributed denial-of-service attack, is an attempt to make a computer resource unavailable to its intended users. Although the means to carry out a DoS attack may vary, it generally consists of a concerted effort to prevent an Internet site or service from functioning efficiently.

One common attack involves saturating the target's (victim's) machine with external communications requests, such that it cannot respond to legitimate traffic, or responds so slowly as to be rendered effectively unavailable. DoS attacks are implemented by either forcing the targeted computer(s) to reset, or consuming its resources so that it can no longer provide its intended service.

The **Denial of Service** screen displays attack type, number of occurrences, and time of last occurrence.

To view the denial of service statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Firewall > Denial of Service** from the left-hand side of the UI.

The screenshot shows the 'Denial of Service' screen in the Wireless Controller interface. The table lists various attack types and their corresponding counts and last occurrences. The 'TCP Bad Sequence' attack type has the highest count at 14, followed by 'Twinge' at 11 and 'IP Spoof' at 2. Most other attack types have a count of 0 and a last occurrence of 'Never'.

Attack Type	Count	Last Occurrence
Ascend	0	Never
Broadcast/Multicast ICMP	0	Never
Chargen	0	Never
Fraggle	0	Never
FTP Bounce	0	Never
Router Solicit	0	Never
Invalid Protocol	0	Never
LAND	0	Never
Router Advertisement	0	Never
Smurf	0	Never
Snork	0	Never
Source Route	0	Never
IP Spoof	2	3 days 09:40:18 ago
TCP Bad Sequence	14	2 days 03:57:50 ago
TCP FIN Scan	0	Never
TCP Header Fragment	0	Never
TCP Intercept	0	Never
TCP IP TTL Zero	0	Never
TCP NULL Scan	0	Never
TCP Post SYN	0	Never
TCP XMAS Scan	0	Never
Twinge	11	4 days 00:02:38 ago
UDP Short Header	0	Never
WINNIKE	0	Never

FIGURE 460 Wireless Controller Firewall DoS screen

4. The **Denial of Service** screen displays the following:

- Attack Type** Displays the DoS attack type. The controller supports enabling or disabling 24 different DoS attack filters.
- Count** Displays the number of times each DoS attack was observed by the controller firewall.
- Last Occurrence** Displays the amount of time since the DoS attack has been observed by the controller firewall.

4. The **IP Firewall Rules** screen displays the following:

Precedence	Displays the precedence value applied to packets. Every rule has a unique precedence value between 1 and 5000. You cannot add two rules with the same precedence value.
Friendly String	This is a string that provides more information as to the contents of the rule. This is for information purposes only.
Hit Count	Displays the number of times each WLAN ACL has been triggered.

MAC Firewall Rules

Firewall

The ability to allow or deny client access by MAC address ensures malicious or unwanted users are unable to bypass security filters. Firewall rules can use one of the three following actions based on a rule criteria:

- *Allow a connection*
- *Allow a connection only if it is secured through the MAC firewall security*
- *Block a connection*

To view MAC firewall rules:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Firewall > MAC Firewall Rules** from the left-hand side of the controller UI.

Precedence	Friendly String	Hit Count
	firewall1	10

FIGURE 462 Wireless Controller MAC Firewall Rules screen

4. The **MAC Firewall Rules** screen displays the following:

Precedence	Displays the precedence value, which are applied to packets. Every rule has a unique precedence value between 1 and 5000. You cannot add two rules with the same precedence value.
Friendly String	This is a string that provides more information as to the contents of the rule. This is for information purposes only.
Hit Count	Displays the number of times each WLAN ACL has been triggered.

NAT Translations

Firewall

1. Select the **Statistics** menu from the Web UI.
2. Select an **Access Point** node from the left navigation pane.
3. Select **Firewall > NAT Translations**.

Protocol	Forward Source IP	Forward Source Port	Forward Dest IP	Forward Dest Port	Reverse Source IP	Reverse Source Port	Reverse Dest IP	Reverse Dest Port
UDP	7.7.7	777	3.3.3	333	8.8.8	888	4.4.4	444

FIGURE 463 Access Point Firewall NAT Translation screen

4. The NAT Translations screen displays the following:

Protocol	Displays the IP protocol type, either UDP or TCP.
Forward Source IP	Displays the internal network IP address for forward facing NAT translations in the Forward Source IP column.
Forward Source Port	Displays the internal network port for forward facing NAT translations in the Forward Source Port column.
Forward Dest IP	Displays the external network destination IP address for forward facing NAT translations in the Forward Dest IP column.
Forward Dest Port	Displays the external network destination port for forward facing NAT translations in the Forward Dest Port column.
Reverse Source IP	Displays the internal network IP address for reverse facing NAT translations in the Reverse Source IP column.
Reverse Source Port	Displays the internal network port for reverse facing NAT translations in the Reverse Source Port column.
Reverse Dest IP	Displays the external network destination IP address for reverse facing NAT translations in the Reverse Dest IP column.
Reverse Dest Port	Displays the external network destination port for reverse facing NAT translations in the Reverse Dest Port column.

Viewing DHCP Snooping Statistics

Firewall

When DHCP servers are allocating IP addresses to the clients, DHCP snooping can strengthen the security on the LAN allowing only clients with specific IP/MAC addresses.

To view the DHCP snooping statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Firewall > DHCP Snooping** from the left-hand side of the controller UI

MAC Address	Node Type	IP Address	Netmask	VLAN	Lease Time	Time Elapsed Since Last Update
22-33-44-55-66-77	Router	4.4.4.4	0	100	5m 0s	10s

FIGURE 464 Wireless Controller Firewall DHCP Snooping screen

4. The **DHCP Snooping** screen displays the following:

MAC Address	Displays the MAC address of the client.
Node Type	Displays the NetBios node with an IP pool from which IP addresses can be issued to client requests on this interface.
IP Address	Displays the IP address used for DHCP discovery and requests between the DHCP server and DHCP clients.
Netmask	Displays the subnet mask used for DHCP discovery and requests between the DHCP server and DHCP clients.

VLAN	Displays the controller interface used for a newly created DHCP configuration.
Lease Time	When a DHCP server allocates an address for a DHCP client, the client is assigned a lease (which expires after a designated interval defined by the administrator). The lease is the time an IP address is reserved for re-connection after its last use. Using short leases, DHCP can dynamically reconfigure networks in which there are more computers than available IP addresses. This is useful, for example, in education and customer environments where client users change frequently. Use longer leases if there are fewer users.
Last Updated	Displays the time the server was last updated.

IPsec

[Wireless Controller Statistics](#)

Use IPsec *Virtual Private Network* (VPN) to define secure tunnels between two peers. Configure sensitive packets, which should be sent through secure tunnels. Once configured, an IPsec peer creates a secure tunnel and sends the packets through the tunnel to the remote peer.

For more information, see:

- [Viewing Security Associations](#)
- [Viewing ISAKMP Statistics](#)

Viewing Security Associations

[IPsec](#)

IPsec tunnels are sets of *security associations* (SA) established between two peers. The security associations define which protocols and algorithms are applied to sensitive packets, and what keying material is used by the two peers.

To view the security associations statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **IPSec > Security Associations** from the left-hand side of the controller UI.

AH IN	AH OUT	Encryption Algorithm	ESP IN	ESP OUT	Local IP Address	MAC Algorithm	Peer	Protocol
1,242	3,445	4586		5,677	124.24.24.2	455	1.2.3.4	455
2,222	2,222	2222		2,222	222.22.22.22	222	2.2.2.2	455
3,333	3,333	3333	33	3,333	133.33.33.3	455	3.3.3.3	333

FIGURE 465 Wireless Controller IP Sec Security Associations screen

4. The **Security Association** screen displays the following:

AH IN	Displays the inbound <i>Authentication Header</i> (AH).
AH OUT	Displays the outbound AH.
Encryption Algorithm	Displays the encryption algorithm used between the peers.
ESP IN	Displays the <i>Security Parameter Index</i> (SPI) in the <i>Encapsulating Security Payload</i> (ESP) inbound header.
ESP OUT	Displays the SPI in the <i>Encapsulating Security Payload</i> (ESP) outbound.
Local IP Address	Displays the IP address of the local peer in an IPsec association.
MAC Algorithm	Displays the algorithm used with the security association.
Peer	Displays the IP address of the remote peer in an IPsec association.
Protocol	Displays the protocol used with the security association between two peers. <i>Internet Key Exchange</i> (IKE) is a key management protocol used with IPsec.

Viewing ISAKMP Statistics

IPsec

Internet Security Association and Key Management Protocol (ISAKMP) is a protocol for establishing security associations and cryptographic keys in an Internet environment. ISAKMP defines the procedures for authenticating a communicating peer, creation and management of security associations, key generation techniques and threat mitigation.

To review ISAKMP stats:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **IPsec > ISAKMP** from the right pane.

Byte Count	Create Time	Encryption Algorithm	Encryption Key Length	Hash Algorithm	Local IP Address	Negotiation Done	Number of Negotiations	Peer	Pseudo Random Function
4,554 bytes	12:45:00	4	4,586 bits	5666	12.2.2.2	✗	566	5.6.7.8	DHh

FIGURE 466 Wireless Controller IPsec ISAKMP screen

4. The **ISAKMP** screen displays the following:

Byte Count	Displays the number of bytes passed between two peers.
Create Time	Displays the exact time when the SA was configured.
Encryption Algorithm	Displays the encryption method used to protect data between peers.
Encryption Key Length	Displays the size (in bits) of the key used in the encryption algorithm.
Hash Algorithm	Defines the hash algorithm used to ensure data integrity. The hash value validates a packet has not been modified in transit.

Local IP Address	Displays the IP address of the local peer in an IPsec association.
Negotiation Done	Displays whether negotiation is completed between peers. During ISAKMP negotiations, peers must identify themselves to one another.
Number of Negotiations	This value is helpful in determining the network address used to validate peers.
Peer	Displays the IP address of the remote peer in an IPsec association.
Pseudo Random Function	Displays the pseudo-random function used to construct keying material for all cryptographic algorithms used by ISAKMP.

Viewing Certificate Statistics

Wireless Controller Statistics

The *Secure Socket Layer* (SSL) protocol is used to ensure secure transactions between Web servers and browsers. This protocol uses a third-party, a certificate authority, to identify one end or both ends of the transactions. A browser checks the certificate issued by the server before establishing a connection.

For more information, see:

- [Viewing Trustpoints Statistics](#)
- [Viewing the RSA Key Details](#)

Viewing Trustpoints Statistics

Viewing Certificate Statistics

Each certificate is digitally signed by a trustpoint. The trustpoint signing the certificate can be a certificate authority, corporate or individual. A trustpoint represents a CA/identity pair containing the identity of the CA, CA-specific configuration parameters and an association with an enrolled identity certificate.

To view the trustpoint statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Certificates > Trustpoints** from the left-hand side of the controller UI.

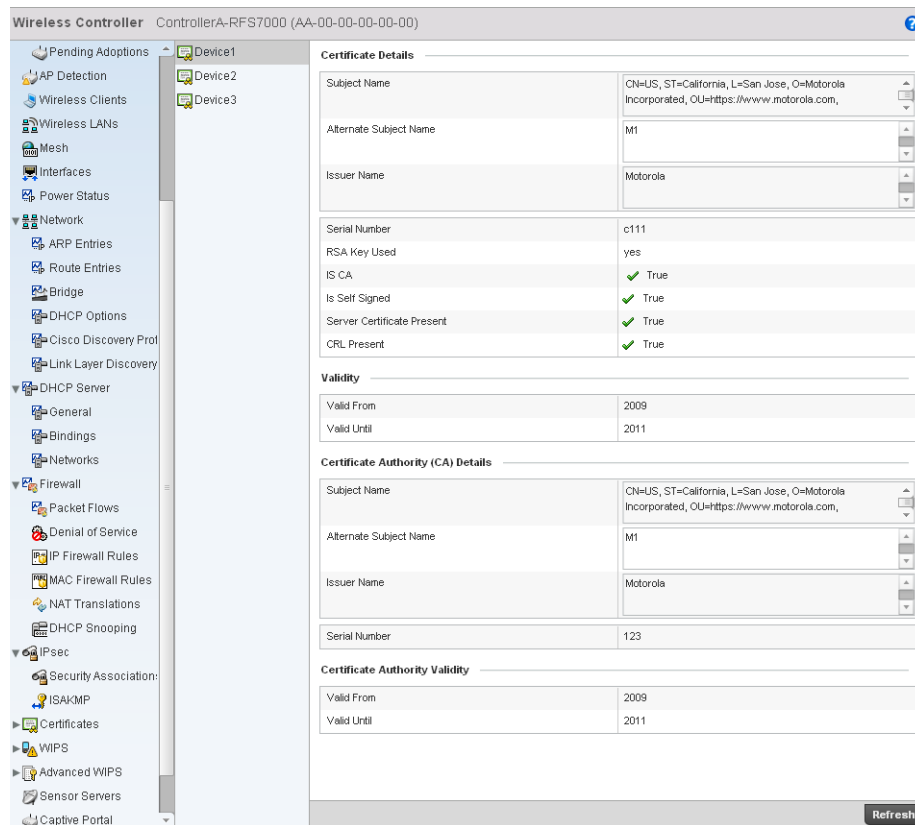


FIGURE 467 Wireless Controller Trustpoints Statistics screen

4. The **Certificate Details** field displays the following:

Subject Name	Describes the entity to which the certificate is issued.
Alternate Subject Name	This section provides alternate information about the certificate as provided to the certificate authority. This field is used to provide information supporting the <i>Subject Name</i> .
Issuer Name	Displays the name of the organization issuing the certificate.
Serial Number	Lists the unique serial number of the certificate.
RSA Key Used	Displays the name of the key pair generated separated, or automatically when selecting a certificate.
IS CA	Indicates if this certificate is an authority certificate.
Is Self Signed	Displays whether the certificate is self-signed. <i>True</i> represents the certificate is self-signed.
Server Certification Present	Displays whether a server certification is present or not. <i>True</i> represents the server certification is present.
CRL Present	Displays whether a <i>Certificate Revocation List</i> (CRL) is present. A CRL contains a list of subscribers paired with digital certificate status. The list displays revoked certificates along with the reasons for revocation. The date of issuance and the entities that issued the certificate are also included.

5. The **Validity** field displays the following:

Valid From	Displays the certificate's issue date.
Valid Until	Displays the certificate's expiration date.

6. The **Certificate Authority (CA) Details** field displays the following:

Subject Name	Displays information about the entity to which the certificate is issued.
Alternate Subject Name	This section provides alternate information about the certificate as provided to the certificate authority. This field is used to provide more information that supports information provided in the <i>Subject Name</i> field.
Issuer Name	Displays the organization issuing the certificate.
Serial Number	The unique serial number of the certificate issued.

7. The **Certificate Authority Validity** field displays the following:

Validity From	Displays the date when the validity of a CA began.
Validity Until	Displays the date when the validity of a CA expires.

Viewing the RSA Key Details

[Viewing Certificate Statistics](#)

Rivest, Shamir, and Adleman (RSA) is an algorithm for public key cryptography. It's the first algorithm known to be suitable for signing as well as encryption.

To view the RSA Key details:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Certificates > RSA Keys** from the left-hand side of the controller UI.

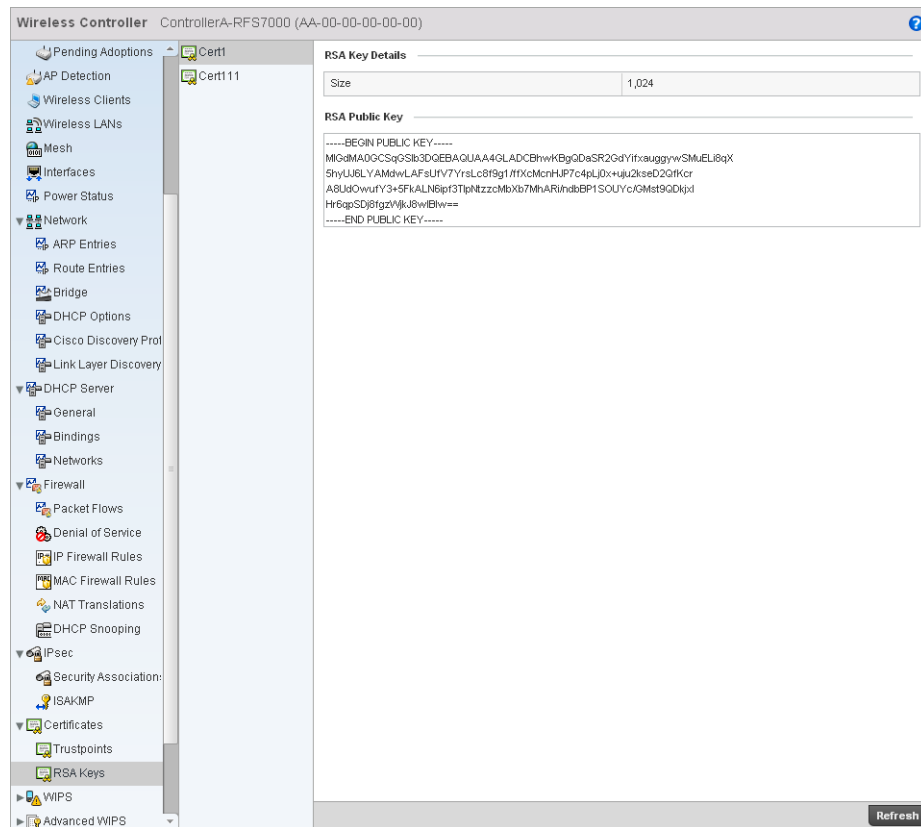


FIGURE 468 Wireless Controller RSA Key Details screen

4. The **RSA Key Details** field describes the size (in bits) of the desired key. If not specified, a default key size of 1024 is used.
5. The **RSA Public Key** field describes the public key used for encrypting messages. This key is known to everyone.

Controller WIPS Statistics

Wireless Controller Statistics

Wireless Intrusion Protection System (WIPS) detects the presence of unauthorized Access Points. Unauthorized attempts to access the WLAN is generally accompanied by intruding clients finding network vulnerabilities. Basic forms of this behavior can be monitored and reported without a dedicated WIPS deployment. When the parameters exceed a configurable threshold, the controller generates a SNMP trap and reports the result via the management interfaces. Basic WIPS functionality does not require monitoring APs and does not perform off-channel scanning.

For more information, see:

- [Viewing Client Blacklist](#)
- [Viewing WIPS Event Statistics](#)

Viewing Client Blacklist

Controller WIPS Statistics

The client blacklist screen displays the statistics for blacklisted clients detected by WIPS.

To view the client blacklist screen:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **WIPS > Client Blacklist** from the left-hand side of the controller UI.

Event Name	Blacklisted Client	Time Blacklisted	Total Time	Time Left
dos-egpol-start-storm	44-55-44-55-44-55	Thu Jun 10 2010 12:26:28 PM	2h 0m 0s	1h 0m 0s
null-probe-response	44-55-44-55-44-55	Thu Jun 10 2010 12:26:28 PM	40m 0s	20m 0s

Row Count: 2

FIGURE 469 Wireless Controller WIPS Client Blacklist screen

4. The **Client Blacklist** screen displays the following:

Event Name	Displays the name of the detected wireless intrusion.
Blacklisted Client	Displays the MAC address of the intruding access point.
Time Blacklisted	Displays the time this client was blacklisted.
Total Time	Displays the length of time the unauthorized device remained in the WLAN.
Time Left	Displays the duration after which the blacklisted client is removed from the blacklist.

Viewing WIPS Event Statistics

To view WIPS event statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **WIPS > WIPS Events** from the left-hand side of the controller UI.

Event Name	Reporting AP	Originating Device	Detector Radio	Time Reported
dos-esp01-start-storm	AP1-ControllerA-AP650	33-44-33-44-33-44	1	Thu Jun 10 2010 12:26:28 PM
null-probe-response	AP1-ControllerA-AP650	33-44-33-44-33-44	1	Thu Jun 10 2010 12:26:28 PM

FIGURE 470 Wireless Controller WIPS Events screen

4. The **WIPS Events** screen displays the following:

Event Name	Displays the name of the detected intrusion event.
Reporting AP	Displays the MAC address of the AP reporting this intrusion.
Originating Device	Displays the MAC address of the intruder AP.
Detector Radio	Displays the type of radio detecting the intrusion.
Time Reported	Displays the time when the intruding AP was detected.

Advanced WIPS

[Wireless Controller Statistics](#)

WIPS monitors for the presence of unauthorized rogue access points and attacks against the managed network.

The *Advanced WIPS* screens supports the following:

- [Viewing General WIPS Statistics](#)
- [Viewing Detected AP Statistics](#)
- [Viewing Detected Clients](#)
- [Viewing Event History](#)

Viewing General WIPS Statistics

Advanced WIPS

The *General WIPS* screen describes WIPS server and sensor address information, version and connection state.

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Advanced WIPS > General** from the left-hand side of the controller UI.

Sensor MAC	Sensor Name	Version	Connected Time	Last Seen Time
00-23-68-22-A9-C4	rfs4000-22A9C4	5.2.65536.4	8/15/2010 10:05:55 PM	8/16/2010 12:58:56 AM
00-23-68-31-19-1C	ap650-31191C	5.2.65536.4	8/15/2010 10:41:55 PM	8/16/2010 12:59:56 AM
bb:11:bb:11:bb:11	Sensor 1	sensor_firmware_1	8/10/2010 12:26:28 PM	8/10/2010 12:26:28 PM

FIGURE 471 Wireless Controller Advanced WIPS screen

4. The **Advanced WIPS Server** field displays the number of ports on the WIPS server.

5. The **Connected Sensors** field displays the following:

- Sensor MAC** Displays the MAC address of each listed sensor AP.
- Sensor Name** Displays the name of each sensor AP.
- Version** Displays each sensor AP's firmware version.
- Connected Time** Displays when the sensor AP connected to the controller.
- Last Seen Time** Displays the number of seconds since the controller last received packets from each sensor AP.

Viewing Detected AP Statistics

Advanced WIPS

The *Detected APs* screen displays network address and connection status for APs within the managed network.

To view detected AP stats:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Advanced WIPS > Detected APs** from the left-hand side of the controller UI.

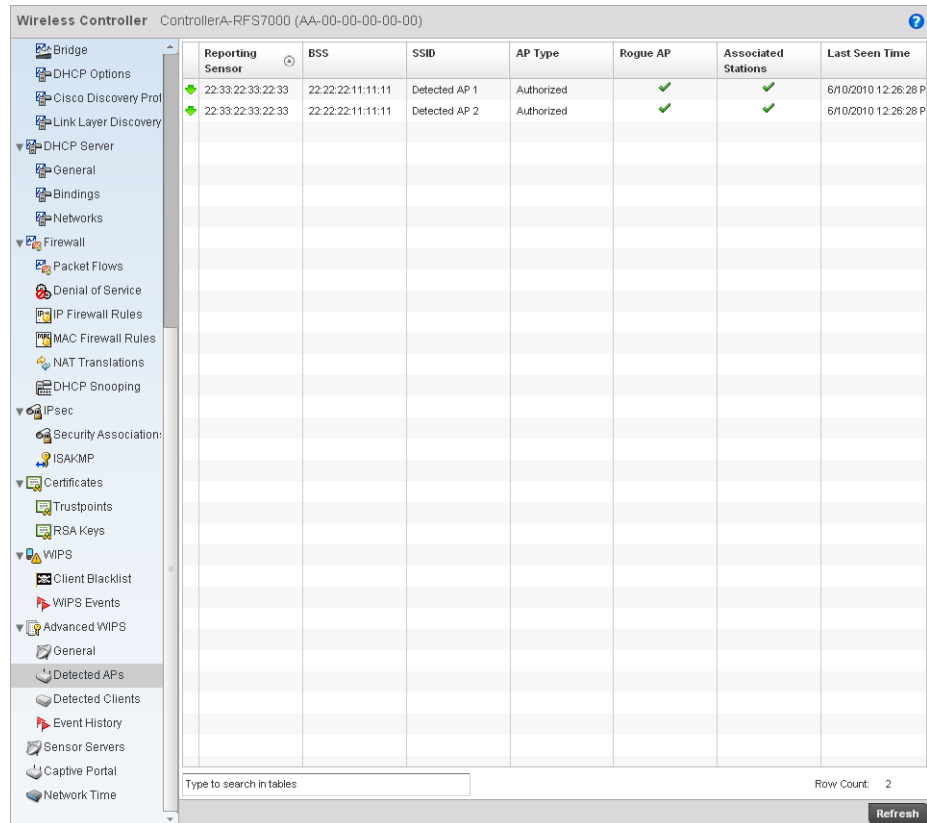


FIGURE 472 Wireless Controller Advanced WIPS Detected APs screen

4. The **Detected APs** screen displays the following:

Reporting Sensor	Displays the numerical value for the radio used with the detecting AP.
BSS	Displays the MAC address of each unapproved AP. These are APs observed on the network, but have yet to be added to the list of approved APs, and are therefore interpreted as a threat.
SSID	Displays the SSID of each unapproved AP. These SSIDs are device SSIDs observed on the network, but have yet to be added to the list of approved APs, and are therefore interpreted as a threat.
AP Type	Displays the type of AP detected.
Rogue AP	Displays the rogue AP's MAC address.
Associated Stations	Displays the number of clients currently associated with the detected AP's radio.
Last Seen Time	Displays the time (in seconds) the unapproved AP was last seen on the network.

Viewing Detected Clients

[Advanced WIPS](#)

The *Detected Clients* screen provides details about rogue clients detected on the network.

To view the detected clients statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Advanced WIPS > Detected Clients** from the left-hand side of the controller UI.

Client MAC	Reporting Sensor	Client Type	Channel	Wired Client	Last Seen Time
55:66:55:66:55:66	77:66:77:66:77:66	Unauthorized	1	✓	6/10/2010 12:26:28 PM

FIGURE 473 Wireless Controller Advanced WIPS Detected Clients screen

4. The **Detected Clients** screen displays the following:

Client MAC	Displays the MAC address of the detected client
Reporting Sensor	Displays the numerical value for the radio used with the detecting AP.
Client Type	Displays the type of client detected.
Channel	Displays the channel the client is transmitting on.
Wired Client	Displays the MAC address detected clients using a wired connection
Last Seen Time	Displays the time (in seconds) the detected client was last seen on the network.

Viewing Event History

Advanced WIPS

The *Event History* screen details unauthorized rogue devices. Unauthorized attempts to access the WLAN are generally accompanied by anomalous behavior, as intruding wireless clients try to find network vulnerabilities.

To view the event history:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.

3. Select **Advanced WIPS > Event History** from the left-hand side of the controller UI.

Event Name	Device MAC	Event Time
wips-event-serious	66:22:66:22:66:22	6/10/2010 12:26:28 PM

FIGURE 474 Wireless Controller Advanced WIPS Event History screen

4. The **Event History** screen displays the following:

Event Name	Displays the name of the detected intrusion.
Device MAC	Displays the MAC address of the intruding device.
Event Time	Displays the time the intruder was detected.

Sensor Server

[Wireless Controller Statistics](#)

Sensor servers allow the monitor and download of data from multiple sensors and remote locations using Ethernet TCP/IP or serial communication. Repeaters are available to extend the transmission range and combine sensors with various frequencies on the same receiver.

To view the Sensor Server statistics:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Sensor Servers** from the left-hand side of the controller UI.

IP Address	Port	Status
1.1.1.1	8,443	connected
2.2.2.2	443	not connected

FIGURE 475 Wireless Controller Sensor Servers screen

4. The **Sensor Servers** screen displays the following:

IP Address	Displays a list of sensor server IP addresses.
Port	Displays the port on which this server is listening.
Status	Displays whether the server is up or down.

Captive Portal Statistics

Wireless Controller Statistics

A captive portal redirects an HTTP client to a Web page (usually for authentication purposes) before authenticating for Internet access. A captive portal turns a Web browser into an authenticator. This is done by intercepting packets (regardless of the address or port) until the user opens a browser and attempts to access the Internet. At that time, the browser is redirected to a Web page requiring authentication.

To view the controller captive portal statistics:

1. Select the **Statistics** tab from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Captive Portal** from the left-hand side of the controller UI.

The screenshot shows the 'Captive Portal' configuration page in the Wireless Controller interface. The page title is 'ControllerA-RFS7000 (AA-00-00-00-00-00)'. On the left is a navigation tree with categories like Bridge, DHCP, Firewall, IPsec, WIPS, and Sensor Servers. The 'Captive Portal' section is selected. The main area displays a table with the following data:

Client MAC	Client IP	Captive Portal	Authentication	WLAN	VLAN	Remaining Time
AA-11-11-00-00-00	1.1.1.1	default	Success	WLAN3	1	1m 40s
AA-11-12-00-00-00	1.1.1.1	default	Pending	WLAN4	2	3m 20s

At the bottom of the table, there is a search input field labeled 'Type to search in tables' and a 'Refresh' button. The 'Row Count' is displayed as 2.

FIGURE 476 Wireless Controller Captive Portal screen

4. The **Captive Portal** screen displays the following:

Client MAC	Displays the requesting client's MAC address.
Client IP	Displays the requesting client's IP address.
Captive Portal	Displays the captive portal page's IP address.
Authentication	Displays the authentication status of the requesting client.
WLAN	Displays the name of the WLAN the client belongs to.
VLAN	Displays the name of the requesting client's VLAN.
Remaining Time	Displays the time after which the client is disconnected from the Internet.

Network Time

[Wireless Controller Statistics](#)

Network Time Protocol (NTP) is central to networks that rely on their wireless controller to supply system time. Without NTP, controller time is unpredictable, which can result in data loss, failed processes, and compromised security. With network speed, memory, and capability increasing at an exponential rate, the accuracy, precision, and synchronization of network time is essential in a controller-managed enterprise network. The wireless controller can use a dedicated server to supply system time. The controller can also use several forms of NTP messaging to sync system time with authenticated network traffic.

Viewing NTP Status

Network Time

The *NTP Status* screen displays performance (status) information relative to the AP's current NTP association. Verify the controller's NTP status to assess the controller's current NTP resource.

To view the NTP status of a managed network:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Network Time > NTP Status** from the left-hand side of the UI.

NTP Status		NTP Association						
Clock Offset	Frequency	Leap	Precision	Reference Time	Reference	Root Delay	Root Dispersion	Stratum
45	11.4	5677	111	dd	344	4566	4899	

FIGURE 477 System NTP screen

4. Refer to the **NTP Status** table to review the accuracy and performance of the controller's synchronization with an NTP server.

Clock Offset	Displays the time differential between the controller time and the NTP resource.
Frequency	An SNTP server clock's skew (difference) for the controller.
Leap	Indicates if a second is added or subtracted to SNTP packet transmissions, or if transmissions are synchronized.
Precision	Displays the precision of the controller's time clock (in Hz). The values that normally appear in this field range from -6 for mains-frequency clocks to -20 for microsecond clocks.
Reference Time	Displays the time stamp the local clock was last set or corrected.
Reference	Displays the address of the time source the controller is synchronized to.
Root Delay	The total round-trip delay in seconds. This variable can take on both positive and negative values, depending on relative time and frequency offsets. The values that normally appear in this field range from negative values (a few milliseconds) to positive values (several hundred milliseconds).
Root Dispersion	The difference between the time on the root NTP server and its reference clock. The reference clock is the clock used by the NTP server to set its own clock.
Status Stratum	Displays how many hops the controller is from its current NTP time source.

Viewing NTP Associations

Network Time

The interaction between the controller and an SNTP server constitutes an association. SNTP associations can be either peer associations (the controller synchronizes to another system or allows another system to synchronize to it), or a server associations (only the controller synchronizes to the SNTP resource, not the other way around).

To view the NTP associations:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Controller** node from the left navigation pane.
3. Select **Network > NTP Association** from the left-hand side of the UI.

NTP Status										
NTP Association										
Delay Time	Display	Offset	Poll	Reach	Reference IP Address	Server IP Address	State	Status	Time	
10	45	67	44	445	12.34.44.44	12.2.2.2	455	ss	now	

FIGURE 478 AP Network Time screen

4. The **NTP Associations** screen provides the controller's current NTP associations.

This screen provides the following:

Delay Time	Displays the round-trip delay (in seconds) for SNTP broadcasts between the SNTP server and the wireless controller.
Dispersion	Displays the time difference between the peer NTP server and the onboard wireless controller clock.
Offset	Displays the calculated offset between the wireless controller and the SNTP server. The controller adjusts its clock to match the server's time value. The offset gravitates towards zero overtime, but never completely reduces its offset to zero.
Poll	Displays the maximum interval between successive messages (in seconds) to the nearest power of two.
Reach	Displays the status of the last eight SNTP messages. If an SNTP packet is lost, the lost packet is tracked over the next eight SNTP messages.
Reference IP Address	Displays the address of the time source the wireless controller is synchronized to.
Server IP Address	Displays the numerical IP address of the SNTP resource (server) providing SNTP updates to the wireless controller.

State	<p>Displays the NTP association status. The state can be one of the following:</p> <ul style="list-style-type: none"> • <i>Synced</i> – Indicates the wireless controller is synchronized to this NTP server. • <i>Unsynced</i> – Indicates the wireless controller has chosen this master for synchronization. However, the master itself is not yet synchronized to UTC. • <i>Selected</i> – Indicates this NTP master server will be considered the next time the wireless controller chooses a master to synchronize with. • <i>Candidate</i> – Indicates this NTP master server may be considered for selection the next time the wireless controller chooses a NTP master server. • <i>Configured</i> – Indicates this NTP server is a configured server.
Stratum	Displays the NTP peer's stratum level.
When	Displays the timestamp of the last NTP packet received from the NTP peer.

Wireless Client Statistics

The *Wireless Client* statistics screen displays read-only statistics for each detected client. It provides an overview of the health of wireless clients in the network. Wireless client statistics includes RF quality, traffic utilization, user details, etc. Use this information to assess if configuration changes are required to improve network performance.

The wireless clients statistics screen can be divided into:

- [Health](#)
- [Details](#)
- [Traffic](#)
- [WMM TSPEC](#)

Health

[Wireless Client Statistics](#)

The **Health** screen displays information on the overall performance of a wireless client.

To view the health of wireless clients:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Client** node from the left navigation pane.
3. Select **Health** from the left-hand side of the controller UI.

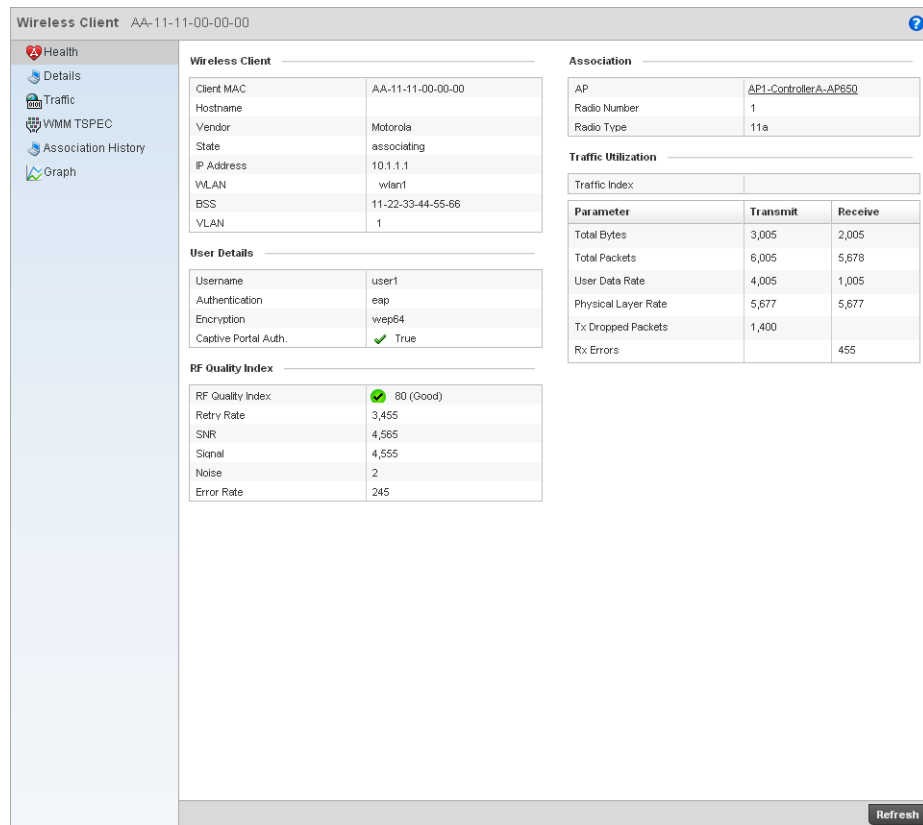


FIGURE 479 Wireless Clients Health screen

4. The **Wireless Client** field displays the following:

Client MAC	Displays the MAC addresses of managed clients.
Vendor	Displays each client's manufacturer.
State	Displays the state of the wireless client. It can be <i>idle</i> , <i>authenticated</i> , <i>associated</i> or <i>blacklisted</i> .
IP Address	Displays the IP address of the wireless client.
WLAN	Displays each client's WLAN name.
BSS	Displays the client's network BSS ID.
VLAN	Displays the client's VLAN ID

5. The **User Details** field displays the following:

Username	Displays the unique administrator or operator name.
Authentication	Lists if authentication is applied. If there's authentication, the status displays.
Encryption	Displays if encryption is applied.
Captive Portal Authentication	Displays whether captive portal authentication is enabled.

6. The **RF Quality Index** field displays the following:

RF Quality Index	Displays client RF quality as a percentage of the connect rate in both directions, as well as the retry and error rate. RF quality index can be interpreted as: <ul style="list-style-type: none"> • 0–20 – very poor quality • 20–40 – poor quality • 40–60 – average quality • 60–100 – good quality
Retry Rate	Displays the average number of retries per packet. A high number indicates possible network or hardware problems.
SNR	Displays the signal to noise ratio of the wireless client associated with the wireless controller.
Signal	Displays radio transmit power in dBm.
Noise	Displays disturbing influences on the signal by interference (in dBm).
Error Rate	Displays the number of received bit rates that have been altered due to noise, interference, and distortion. It is a unitless performance measure.

7. The **Association** field displays the following:

AP	Displays the wireless client's associated AP. Select an AP to view information in detail.
Radio Number	Displays the AP radio the client is associated with.
Radio Type	Displays the radio type as either 802.11b, 802.11bg, 802.11bgn, 802.11a or 802.11an.

8. The **Traffic Utilization** field displays statistics on the traffic generated and received by this wireless client. This area displays the traffic index, which measures how efficiently the traffic medium is used. It is defined as the percentage of current throughput relative to the maximum possible throughput.

Traffic indices are:

- 0–20 – very low utilization
- 20–40 – low utilization
- 40–60 – moderate utilization
- 60 and above – high utilization

This field also displays the following:

Total Bytes	Displays the total bytes processed by the wireless client.
Total Packets	Displays the total number of packets processed by the wireless client.
User Data Rate	Displays the average user data rate.
Physical Layer Rate	Displays the average packet rate at the physical layer.
Tx Dropped Packets	Displays the number of packets dropped during transmission.
Rx Errors	Displays the number of errors encountered during data transmission. The higher the error rate, the less reliable the connection or data transfer.

Details

Wireless Client Statistics

The *Details* screen provides information on a selected wireless client.

To view the details screen of a wireless client:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Client** node from the left navigation pane.
3. Select **Details** from the left-hand side of the controller UI.

The screenshot displays the 'Wireless Client' details page for client ID AA-11-11-00-00-00. The interface is divided into several sections:

- Wireless Client:**
 - SSID: wlan1
 - Hostname: 11-aa-bb-cc-dd-ee
 - RF Domain: RF_Domain1
- User Details:**
 - Username: user1
 - Authentication: eap
 - Encryption: wsep64
 - Captive Portal Auth.: True
- Connection:**
 - Idle Time: 5m 44s
 - Last Active: 10
 - Last Association: 7m 35s
 - Session Time: 7m 35s
 - SM PowerSave Mode: true
 - Power Save Mode: False
 - WMM Support: True
 - 40 MHz Capable: True
 - Max. Physical Rate: 100,000
 - Max. User Rate: 50,000
 - MC2UC Streams: (empty)
- Association:**
 - AP: 11-aa-bb-cc-dd-ee
 - BSS: 11-22-33-44-55-66
 - Radio Number: 1
 - Radio Type: 11bg
 - Rate: 345
- 802.11 Protocol:**
 - High-Throughput: Supported
 - RIFS: Unsupported
 - Unscheduled APSD: 33
 - AID: 22
 - Max. AMSDU Size: 234
 - Max. AMPDU Size: 233
 - Interframe Spacing: 233
 - Short Guard Interval: Unsupported

A 'Refresh' button is located at the bottom right of the screen.

FIGURE 480 Wireless Clients Details screen

4. The **Wireless Client** area displays the following:

SSID Displays the client's associated SSID.

RF Domain Displays the wireless client's RF Domain.

5. The **User Details** field displays the following:

Username Displays the administrator or operator name.

Authentication	Displays whether authentication is invoked. If authentication is applied, the field displays its status.
Encryption	Displays if any encryption is applied.
Captive Portal Auth.	Displays whether captive portal authentication is enabled.

6. The **Connection** field displays the following:

Idle Time	Displays the wireless client's idle time.
Last Active	Displays (in seconds) when the wireless client last communicated with its connected AP.
Last Association	Displays the client's association duration.
Session Time	Displays the duration for which a session can be maintained by the wireless client without it being dis-associated from the system.
SM Power Save Mode	Displays whether SM Power Save is enabled on the wireless client. The <i>spatial multiplexing</i> (SM) power save mode allows an 802.11n client to power down all but one of its radios. This power save mode has two sub modes of operation: static operation and dynamic operation.
Power Save Mode	Displays whether this feature is enabled or not. To prolong battery life, the 802.11 standard defines an optional Power Save Mode, which is available on most 802.11 NICs. End users can simply turn it on or off via the card driver or configuration tool. With power save off, the 802.11 network card is generally in receive mode listening for packets and occasionally in transmit mode when sending packets. These modes require the 802.11 NIC to keep most circuits powered-up and ready for operation.
WMM Support	Displays whether WMM support is enabled.
40 MHz Capable	Displays whether the wireless client has 802.11n channel support operating at 40 MHz.
Max Physical Rate	Displays the maximum data rate at the physical layer.
Max User Rate	Displays the maximum permitted user data rate.

7. The **Association** field displays the following:

AP	Displays the MAC address of the client's associated AP.
BSS	Displays the basic service set the AP belongs to. A BSS is a set stations that can communicate with one another.
Radio Number	Displays the AP radio the wireless client is associated with.
Radio Type	Displays the radio type as either 802.11b, 802.11bg, 802.11bgn, 802.11a or 802.11an.
Rate	Displays the permitted data rate.

8. The **802.11 Protocol** field displays the following:

High-Throughput	Displays whether this feature is supported or not. High throughput is a measure of the successful packet delivery over a communication channel.
RIFS	Displays whether this feature is supported. RIFS is a required 802.11n feature that improves performance by reducing the amount of dead time between OFDM transmissions.
Unscheduled APSD	Displays whether an unscheduled service period is supported as a contiguous period the controller is expected to be awake.
AID	Displays the Association ID established by an AP. 802.11 association enables the access point to allocate resources and synchronize with a radio NIC. An NIC begins the association process by sending an association request to an access point. This association request is sent as a frame. This frame carries information about the NIC and the SSID of the network it wishes to associate. After receiving the request, the access point considers associating with the NIC, and reserves memory space for establishing an AID for the NIC.
Max AMSDU Size	Displays the maximum AMSDU size. AMSDU is a set of Ethernet frames wrapped in a 802.11n frame. This value is the maximum AMSDU frame size in bytes.
Interframe Spacing	Displays the interval between two consecutive Ethernet frames.
Short Guard Interval	Displays the guard interval in micro seconds. Guard intervals prevent interference between distinct data transmissions while.

Traffic

[Wireless Client Statistics](#)

The traffic screen provides an overview of client traffic utilization. This screen also displays a RF quality index.

To view the traffic statistics of a wireless clients:

1. Select the **Statistics** menu from the Web UI.
2. Select a **Wireless Client** node from the left navigation pane.
3. Select **Traffic** from the left-hand side of the controller UI.

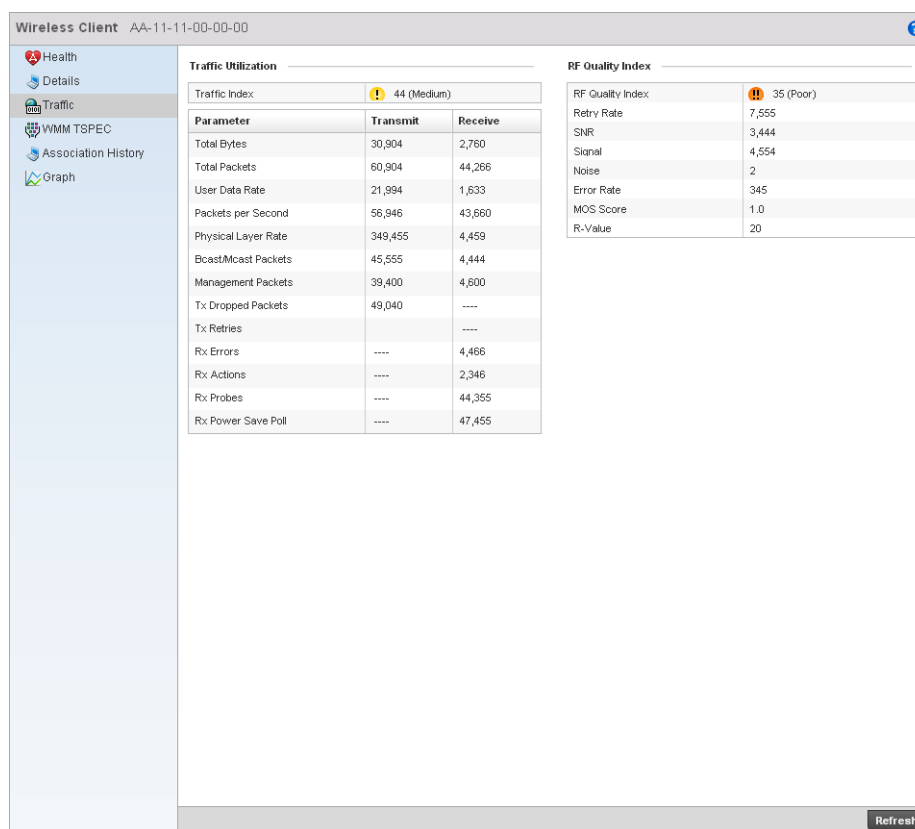


FIGURE 481 Wireless Clients Traffic screen

4. The **Traffic Utilization** field provides the traffic index, which measures how efficiently the traffic medium is used. It is defined as the percentage of current throughput relative to the maximum possible throughput.

Total Bytes	Displays the total bytes processed by the client.
Total Packets	Displays the total number of data packets processed by the wireless client.
User Data Rate	Displays the average user data rate.
Packets per Second	Displays the packets processed per second.
Physical Layer Rate	Displays the data rate at the physical layer level.
Bcast/Mcast Packets	Displays the total number of broadcast/management packets processed.
Management Packets	Displays the number of management packets processed.
Tx Dropped Packets	Displays the number of dropped packets while transmitting.
Tx Retries	Displays the total number of transmit retries.
Rx Errors	Displays the degree of errors encountered during data transmission. The higher the error rate, the less reliable the connection or data transfer.

Rx Actions	Displays the number of receive actions during data transmission.
Rx Probes	Displays the number of probes sent. A probe is a program or other device inserted at a key juncture in a for network for the purpose of monitoring or collecting data about network activity.
Rx Power Save Poll	Displays the power save using the Power Save Poll (PSP) mode. <i>Power Save Poll</i> (PSP) is a protocol, which helps to reduce the amount of time a radio needs to powered. PSP allows the WiFi adapter to notify the access point when the radio is powered down. The access point holds any network packet to be sent to this radio.

5. The **RF Quality Index** field displays the following:

RF Quality Index	Displays the client's RF quality. The RF quality index is the overall effectiveness of the RF environment, as a percentage of the connect rate in both directions as well as the retry rate and the error rate. RF quality index value can be interpreted as: <ul style="list-style-type: none"> • 0-20 – very poor quality • 20-40 –poor quality • 40-60 –average quality • 60-100 – good quality
Retry Rate	Displays the average number of retries per packet. A high number indicates possible network or hardware problems.
SNR	Displays the signal to noise ratio of the wireless client associated with the wireless controller.
Signal	Displays the power of the radio signals in dBm.
Noise	Displays disturbing influences on the signal by interference.
Error Rate	Displays the number of received bit rates altered due to noise, interference, and distortion. It's a unitless performance measure.
MOS Score	Displays the average call quality using the <i>Mean Opinion Score</i> (MOS) call quality scale. The MOS scale rates call quality on a scale of 1-5, with higher scores being better. If the MOS score is lower than 3.5, it's likely users will not be satisfied with the voice quality.
R-Value	Displays the R-value. R-value is a number or score that is used to quantitatively express the quality of speech in communications systems. This is used in digital networks that carry <i>Voice over IP</i> (VoIP) traffic. The R-value can range from 1 (worst) to 100 (best) and is based on the percentage of users who are satisfied with the quality of a test voice signal after it has passed through a network from a source (transmitter) to a destination (receiver). The R-value scoring method accurately portrays the effects of packet loss and delays in digital networks carrying voice signals.

WMM TSPEC

[Wireless Client Statistics](#)

802.11e *Traffic Specification* (TSPEC) provides a set of parameters that defines the characteristics of the traffic stream, for example: operating requirement and scheduling. The sender T spec specifies parameters available for the flow of packets. Both the sender and the receiver use Tspec.

The TSPEC screen provides the information about these TSPEC counts, TSPEC types, etc. of the wireless client selected.

To view the TSPEC statistics:

1. Select the **Statistics** from the Web UI.

2. Select a **Wireless Client** node from the left navigation pane.
3. Select **WMM TSPEC** from the left-hand side of the controller UI.

FIGURE 482 Wireless Clients 802.11e TSPEC screen

4. The **TSPEC Count** displays the number of TSPECs available for the client's packet flow.
5. The **TSPEC Type** field displays the following:

Voice	Displays the status of voice traffic prioritization. A red 'X' indicates this feature is disabled. A green check mark indicates this feature is enabled.
Video	Displays the status of prioritization for video traffic. A red 'X' indicates this feature is disabled. A green check mark indicates this feature is enabled.
Best Effort	Displays the status of prioritization for best effort traffic. A red 'X' indicates this feature is disabled. A green check mark indicates this feature is enabled.
Background	Displays the status of prioritization for background traffic. A red 'X' indicates this feature is disabled. A green check mark indicates this feature is enabled.

6. The **First TSPEC** field displays the following:

Direction	Displays whether this is a transmitting or receiving flow.
Parameter	Displays the parameter for defining the traffic stream. <i>TID</i> identifies data packets as belonging to a unique traffic stream.
Voice	Displays the voice corresponding to TID and Media Time.

Video	Displays the Video corresponding to TID and Media Time.
Best Effort	Displays the Best Effort corresponding to TID and Media Time.
Background	Displays the Background corresponding to TID and Media Time.

7. The **Second TSPEC** field displays the following:

Direction	Displays whether this is a transmitting or receiving flow
Parameter	Displays the parameter for defining the traffic stream. <i>TID</i> identifies data packets as belonging to a unique traffic stream.
Voice	Displays the voice corresponding to TID and Media Time.
Video	Displays the Video corresponding to TID and Media Time.
Best Effort	Displays the Best Effort corresponding to TID and Media Time.
Background	Displays the Background corresponding to TID and Media Time.

Publicly Available Software

In this appendix

- [General Information](#) 723
- [Open Source Software Used](#) 724
- [OSS Licenses](#) 729

General Information

This document contains information regarding licenses, acknowledgments and required copyright notices for open source packages used in this Brocade product.

A

Open Source Software Used

Wireless Controller

Name	Version	Origin	License
Linux kernel	2.6.16.51	http://www.kernel.org	gplv2
bridge-utils	1.0.4	http://www.kernel.org	gplv2
pciutils	2.1.11 & 2.1.11-15.patch	http://mj.ucw.cz/pciutils.html	gplv2
busybox	1.1.3	http://www.busybox.net	gplv2
LILO	22.6	http://lilo.go.dyndns.org	bsd
e2fsprogs	busybox-1.1.3	http://e2fsprogs.sourceforge.net/	gplv2
Zlib	1.2.3	http://www.zlib.net	bsd
ethereal	0.10.14	http://www.ethereal.com	gplv2
strace	4.5.19	http://sourceforge.net/projects/strace/	bsd and gplv2
glibc	2.7	http://www.gnu.org	lgplv2
glib2	2.7.0	http://www.gtk.org/	gplv2
gdb	6.5	http://www.gnu.org	gplv2
safestr	1.0.3	http://www.zork.org/safestr	bsd
iproute2	50816	http://developer.osdl.org	gplv2
iptables	1.3.5	http://www.netfilter.org/	gplv2
libdnet	1.1	http://libdnet.sourceforge.net/	bsd
libncurses	5.4	http://www.gnu.org/software/ncurses/ncurses.html	MIT
libpcap	0.9.4	http://www.tcpdump.org/	bsd
tcpdump	3.9.7	http://www.tcpdump.org/	bsd
libreadline	4.3	http://tiswww.case.edu/php/chet/readline/rltop.html	gplv2
ntp	4.2.4p8	http://www.ntp.org/	bsd
mii-diag	2.09	http://www.scyld.com	gplv2
mtd-tools	1.3.1 (mtd-utils-1.3.1?)	http://www.linux-mtd.infradead.org	gplv2
DHCP	3.0.3	http://www.isc.org	bsd
MD5 hashing	BusyBox v1.1.3	http://www.ietf.org/rfc/rfc1321.txt	bsd
Kerberos client	5	http://www.mit.edu/~kerberos	bsd
Authentication modules		http://www.kernel.org/pub/linux/libs/pam/	gplv2
diff utility	2.8.1	http://www.gnu.org/software/diffutils/diffutils.html	gplv2
nano editor	1.2.4	http://www.nano-editor.org	gplv2

Name	Version	Origin	License
thttpd	2.25b	http://www.acme.com	bsd
net-snmp	5.3.0.1	http://net-snmp.sourceforge.net	bsd
smidump library	0.4.3	http://www.ibr.cs.tu-bs.de/projects/libsmi/index.html	bsd
OpenSSH	5.4p1	http://www.openssh.com	bsd
OpenSSL	0.9.8n	http://www.openssl.org	openssl
stunnel	4.31	http://www.stunnel.org	gplv2
qdbm	1.8.77	http://qdbm.sourceforge.net/	lgplv2
advas	0.2.3	http://sourceforge.net/projects/advas/	gplv2
libexpat	2.0.0	http://expat.sourceforge.net/	mit
ppp	2.4.4	http://ppp.samba.org/	bsd
openldap	2.3.20	http://www.openldap.org/	bsd
pure-ftpd	1.0.22	http://www.pureftpd.org	bsd
FreeRADIUS	2.1.7	http://freeradius.org/	gplv2
rp-pppoe	3.1	http://www.roaringpenguin.com/products/pppoe	gplv2
Stackless python	252	http://www.stackless.com/	bsd
xxl	1.0.1	http://zork.org/xxl/	bsd
libgmp	4.2.2	http://gmplib.org/	lgplv2
procname	0.2	http://code.google.com/p/procname/	lgplv2
bind	9.4.3-P3	https://www.isc.org/software/bind	bsd
dosfstools	2.11	http://www.daniel-baumann.ch/software/dosfstools/	gplv2
ebtables	v2.0.6	http://ebtables.sourceforge.net/	gplv2
procps	3.2.7	http://procps.sourceforge.net/	gplv2
libxml	2.7.3	http://xmlsoft.org/	MIT
libpopt	1.14.4	http://packages.debian.org/changelogs/pool/main/p/popt/	MIT
libusb	0.1.12	http://www.libusb.org/	lgplv2
sysstat	9.0.3	http://sebastien.godard.pagesperso-orange.fr/	gplv2
pychecker	0.8.18	http://pychecker.sourceforge.net/	bsd

AP650

Name	Version	Origin	License
autoconf	2.62	http://www.gnu.org/software/autoconf/	gplv2
automake	1.9.6	http://www.gnu.org/software/automake/	gplv2
binutils	2.19.1	http://www.gnu.org/software/binutils/	gplv2
bison	2.3	http://www.gnu.org/software/bison/	gplv2

A

Name	Version	Origin	License
busybox	1.11.3	http://www.busybox.net/	gplv2
dnsmasq	2.47	http://www.thekelleys.org.uk/dnsmasq/doc.html	gplv2
dropbear	0.51	http://matt.ucc.asn.au/dropbear/dropbear.html	dropbear
e2fsprogs	1.40.11	http://e2fsprogs.sourceforge.net/	gplv2
gcc	4.1.2	http://gcc.gnu.org/	gplv2
gdb	6.8	http://www.gnu.org/software/gdb/	gplv2
genext2fs	1.4.1	http://genext2fs.sourceforge.net/	gplv2
glibc	2.7	http://www.gnu.org/software/libc/	gplv2
hostapd	0.6.9	http://hostap.epitest.fi/hostapd/	gplv2
hotplug2	0.9	http://istev.e.bofh.cz/~istev/hotplug2/	gplv2
ipkg-utils	1.7	http://www.handhelds.org/sources.html	gplv2
iproute2	2.6.25	http://www.linuxfoundation.org/collaborate/workgroups/networking/iproute2	gplv2
iptables	1.4.1.1	http://www.netfilter.org/	gplv2
libpcap	0.9.8	http://www.tcpdump.org/	bsd
libtool	1.5.24	http://www.gnu.org/software/libtool/	gplv2
linux	2.6.28.9	http://www.kernel.org/	gplv2
lzma	4.32	http://www.7-zip.org/sdk.html	lgplv2
lzo	2.03	http://www.oberhumer.com/opensource/lzo/	gplv2
m4	1.4.5	http://www.gnu.org/software/m4/	gplv2
madwifi	trunk-r3314	http://madwifi-project.org/	bsd
mtdev	5/5/2009	http://www.linux-mtd.infradead.org/	gplv2
mtdev-utils	2/27/2009	http://www.linux-mtd.infradead.org/	gplv2
openssl	0.9.8j	http://www.openssl.org/	openssl
openwrt	trunk-r15025	http://www.openwrt.org/	gplv2
opkg	trunk-r4564	http://code.google.com/p/opkg/	gplv2
pkg-config	0.22	http://pkg-config.freedesktop.org/wiki/	gplv2
ppp	2.4.3	http://ppp.samba.org/ppp/	bsd
quilt	0.47	http://savannah.nongnu.org/projects/quilt/	gplv2
sed	4.1.2	http://www.gnu.org/software/sed/	gplv2
squashfs	3	http://squashfs.sourceforge.net/	gplv2
u-boot	trunk-2010-03-30	http://www.denx.de/wiki/U-Boot/	gplv2

AP51xx

Name	Version	Origin	License
Linux	2.4.20_mv131-ixdp 4xx	www.mvista.com	gplv2 and MontaVista
Apache Web server	1.3.41	www.apache.org	apache
Java Development Kit libraries	1.5.0_01	http://java.sun.com/j2se/	Sun Community Source
Kerberos Client	5	http://www.mit.edu/~kerberos	bsd
AES/CCM encryption		http://www.gladman.me.uk/	bsd
zlib	1.1.4	http://www.zlib.net/	zlib
freeradius	1.0.0-pre3	http://www.freeradius.org/	gplv2
net-snmp	5.0.9	http://net-snmp.sourceforge.net	bsd
openssh	5.4p1	http://www.openssh.com/	bsd
openldap	2.2	http://www.openldap.org/foundation/	openldap
wuftp	2.6.1	http://wu-ftp.d.therockgarden.ca/	wuftp

AP7131

Name	Version	Origin	License
Apache Web Server	1.3.41	http://www.apache.org/	apache
autoconf	2.62	http://www.gnu.org/software/autoconf/	gplv2
automake	1.9.6	http://www.gnu.org/software/automake/	gplv2
bind	9.3.2	http://www.isc.org/	bsd
binutils	2.19.1	http://www.gnu.org/software/binutils/	gplv2
bison	2.3	http://www.gnu.org/software/bison/	gplv2
bridge	1.0.4	http://www.linuxfoundation.org/collaborate/workgroups/networking/bridge/	gplv2
busybox	1.11.3	http://www.busybox.net/	gplv2
e2fsprogs	1.40.11	http://e2fsprogs.sourceforge.net/	gplv2
flex	2.5.4	http://flex.sourceforge.net/	bsd
freeradius	2.0.2	http://www.freeradius.org/	gplv2
gcc	4.1.2	http://gcc.gnu.org/	gplv2
gdb	6.8	http://www.gnu.org/software/gdb/	gplv2
genext2fs	1.4.1	http://genext2fs.sourceforge.net/	gplv2
glibc	2.7	http://www.gnu.org/software/libc/	gplv2
ipkg-utils	1.7	http://www.handhelds.org/sources.html	gplv2

A

Name	Version	Origin	License
iptables	1.4.3	http://www.netfilter.org/projects/iptables/index.html	gplv2
iproute2	2.6.25	http://www.linuxfoundation.org/collaborate/workgroups/networking/iproute2	gplv2
iptables	1.4.1.1	http://www.netfilter.org/	gplv2
kerberos	5	http://web.mit.edu/Kerberos/	gplv2
libpam	0.99.9.0	http://www.kernel.org/pub/linux/libs/pam/	gplv2
libpcap	0.9.8	http://www.tcpdump.org/	bsd
libtool	1.5.24	http://www.gnu.org/software/libtool/	gplv2
linux	2.6.28.9	http://www.kernel.org/	gplv2
lzma	4.32	http://www.7-zip.org/sdk.html	lgplv2
lzo	2.03	http://www.oberhumer.com/opensource/lzo/	gplv2
mod_ssl	2.8.3.1-1.3.41	http://www.modssl.org/	bsd
mtdev	5/5/2009	http://www.linux-mtd.infradead.org/	gplv2
mtdev-utils	2/27/2009	http://www.linux-mtd.infradead.org/	gplv2
openldap	2.3.20	http://www.openldap.org/foundation/	openldap
openldap	0.0.3alpha	http://openldap.sourceforge.net/	bsd
openssh	5.4p1	http://www.openssh.com/	bsd
openssl	0.9.8j	http://www.openssl.org/	openssl
ppp	2.4.3	http://ppp.samba.org/ppp/	bsd
snmpagent	5.0.9	http://sourceforge.net/	bsd
strace	4.5.18	http://sourceforge.net/projects/strace/	bsd
u-boot	Trunk-2010-03-30	http://www.denx.de/wiki/U-Boot/	gplv2
wireless_tools	r29	http://www.hpl.hp.com/personal/Jean_Tourrilhes/Linux/Tools.html	gplv2
wuftp	1.0.21	http://wu-ftp.therockgarden.ca/	wuftp
zlib	1.2.3	http://www.zlib.net/	zlib

OSS Licenses

GNU General Public License 2.0

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

A

This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
 - a. The modified work must itself be a software library.
 - b. You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
 - c. You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
 - d. If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

A

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

- a. Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)
- b. Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- c. Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- d. If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- e. Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:
 - a. Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.

b. Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.
10. Each time you redistribute the Library (or any work based on the library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.
11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

A

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

GNU Lesser General Public License 2.1

GNU LESSER GENERAL PUBLIC LICENSE

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.

51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages—typically libraries—of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is

modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

A

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
 - a. The modified work must itself be a software library.
 - b. You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
 - c. You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
 - d. If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

A

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

- a. Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

- b. Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- c. Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- d. If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- e. Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

- 7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:
 - a. Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.
 - b. Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.
- 8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
- 9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

A

10. Each time you redistribute the Library (or any work based on the library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.
11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

BSD Style Licenses

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, and the entire permission notice in its entirety, including the disclaimer of warranties.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, ALL OF WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF NOT ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Copyright (c) 1988 Regents of the University of California. All rights reserved. Redistribution and use in source and binary forms are permitted provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed by the University of California, Berkeley. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED ``AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

MIT License

Copyright 1987, 1988 by MIT Student Information Processing Board. Permission to use, copy, modify, and distribute this software and its documentation for any purpose is hereby granted, provided that the names of M.I.T. and the M.I.T. S.I.P.B. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. M.I.T. and the M.I.T. S.I.P.B. make no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

WU-FTPD License

Use, modification, or redistribution (including distribution of any modified or derived work) in any form, or on any medium, is permitted only if all the following conditions are met:

1. Redistributions qualify as "freeware" or "Open Source Software" under the following terms:
 - a. Redistributions are made at no charge beyond the reasonable cost of materials and delivery. Where redistribution of this software is as part of a larger package or combined work, this restriction applies only to the costs of materials and delivery of this software, not to any other costs associated with the larger package or combined work.
 - b. Redistributions are accompanied by a copy of the Source Code or by an irrevocable offer to provide a copy of the Source Code for up to three years at the cost of materials and delivery. Such redistributions must allow further use, modification, and redistribution of the Source Code under substantially the same terms as this license. For the purposes of redistribution "Source Code" means all files included in the original distribution, including all modifications or additions, on a medium and in a form allowing fully working executable programs to be produced.
2. Redistributions of Source Code must retain the copyright notices as they appear in each Source Code file and the COPYRIGHT file, these license terms, and the disclaimer/limitation of liability set forth as paragraph 6 below.
3. Redistributions in binary form must reproduce the Copyright Notice, these license terms, and the disclaimer/limitation of liability set forth as paragraph 6 below, in the documentation and/or other materials provided with the distribution. For the purposes of binary distribution the "Copyright Notice" refers to the following language:

Copyright (c) 1999,2000,2001 WU-FTPD Development Group. All rights reserved.
 Portions Copyright (c) 1980, 1985, 1988, 1989, 1990, 1991, 1993, 1994 The Regents of the University of California. Portions Copyright (c) 1993, 1994 Washington University in Saint Louis. Portions Copyright (c) 1996, 1998 Berkeley Software Design, Inc. Portions Copyright (c) 1998 Sendmail, Inc. Portions Copyright (c) 1983, 1995, 1996, 1997 Eric P. Allman. Portions Copyright (c) 1989 Massachusetts Institute of Technology. Portions Copyright (c) 1997 Stan Barber. Portions Copyright (c) 1991, 1992, 1993, 1994, 1995, 1996, 1997 Free Software Foundation, Inc. Portions Copyright (c) 1997 Kent Landfield.

Use and distribution of this software and its source code are governed by the terms and conditions of the WU-FTPD Software License ("LICENSE").

If you did not receive a copy of the license, it may be obtained online at <http://www.wu-ftp.org/license.html>

4. All advertising materials mentioning features or use of this software must display the following acknowledgement: "This product includes software developed by the WU-FTPD Development Group, the Washington University at Saint Louis, Berkeley Software Design, Inc., and their contributors."
5. Neither the name of the WU-FTPD Development Group, nor the names of any copyright holders, nor the names of any contributors may be used to endorse or promote products derived from this software without specific prior written permission. The names "wuftpd" and "wu-ftpd" are trademarks of the WU-FTPD Development Group and the Washington University at Saint Louis.
6. Disclaimer/Limitation of Liability:

THIS SOFTWARE IS PROVIDED BY THE WU-FTPD DEVELOPMENT GROUP, THE COPYRIGHT HOLDERS, AND CONTRIBUTORS, "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE WU-FTPD DEVELOPMENT GROUP, THE COPYRIGHT HOLDERS, OR CONTRIBUTORS, BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.
7. USE, MODIFICATION, OR REDISTRIBUTION, OF THIS SOFTWARE IMPLIES ACCEPTANCE OF ALL TERMS AND CONDITIONS OF THIS LICENSE.

Open SSL License

LICENSE ISSUES

=====

The OpenSSL toolkit stays under a dual license, i.e. both the conditions of the OpenSSL License and the original SSLeay license apply to the toolkit. See below for the actual license texts. Actually both licenses are BSD-style Open Source licenses. In case of any license issues related to OpenSSL please contact openssl-core@openssl.org.

OpenSSL License

=====

Copyright (c) 1998-2008 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in documentation and/or other materials provided with the distribution.

A

3. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"
4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License

Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com) All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, hash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. All advertising materials mentioning features or use of this software must display the following acknowledgement: "This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)" The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related.
4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement: "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The licence and distribution terms for any publicly available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution licence [including the GNU Public Licence.]

ZLIB License

Copyright (C) 1995-2005 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly Mark Adler
jloup@gzip.org madler@alumni.caltech.edu

Open LDAP Public License

The OpenLDAP Public License Version 2.8, 17 August 2003 Redistribution and use of this software and associated documentation ("Software"), with or without modification, are permitted provided that the following conditions are met:

1. Redistributions in source form must retain copyright statements and notices.
2. Redistributions in binary form must reproduce applicable copyright statements and notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution, and

A

3. Redistributions must contain a verbatim copy of this document.

The OpenLDAP Foundation may revise this license from time to time. Each revision is distinguished by a version number. You may use this Software under terms of this license revision or under the terms of any subsequent revision of the license.

THIS SOFTWARE IS PROVIDED BY THE OPENLDAP FOUNDATION AND ITS CONTRIBUTORS ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OPENLDAP FOUNDATION, ITS CONTRIBUTORS, OR THE AUTHOR(S) OR OWNER(S) OF THE SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The names of the authors and copyright holders must not be used in advertising or otherwise to promote the sale, use or other dealing in this Software without specific, written prior permission.

Title to copyright in this Software shall at all times remain with copyright holders. OpenLDAP is a registered trademark of the OpenLDAP Foundation. Copyright 1999-2003 The OpenLDAP Foundation, Redwood City, California, USA. All Rights Reserved. Permission to copy and distribute verbatim copies of this document is granted.

Apache License 2.0

Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. **Grant of Copyright License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. **Grant of Patent License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
4. **Redistribution.** You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
 - a. You must give any other recipients of the Work or Derivative Works a copy of this License; and
 - b. You must cause any modified files to carry prominent notices stating that You changed the files; and
 - c. You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

d. If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. **Submission of Contributions.** Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. **Trademarks.** This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. **Disclaimer of Warranty.** Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
8. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. **Accepting Warranty or Additional Liability.** While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf

and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Drop Bear License

Dropbear contains a number of components from different sources, hence there are a few licenses and authors involved. All licenses are fairly non-restrictive.

The majority of code is written by Matt Johnston, under the license below.

Portions of the client-mode work are (c) 2004 Mihnea Stoenescu, under the same license:

Copyright (c) 2002-2006 Matt Johnston

Portions copyright (c) 2004 Mihnea Stoenescu

All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

LibTomCrypt and LibTomMath are written by Tom St Denis, and are Public Domain.

A

=====

sshpty.c is taken from OpenSSH 3.5p1,

Copyright (c) 1995 Tatu Ylonen , Espoo, Finland

All rights reserved

"As far as I am concerned, the code I have written for this software can be used freely for any purpose. Any derived versions of this software must be clearly marked as such, and if the derived work is incompatible with the protocol description in the RFC file, it must be called by a name other than "ssh" or "Secure Shell". "

=====

loginrec.c

loginrec.h

atomicio.h

atomicio.c

and strcat() (included in util.c) are from OpenSSH 3.6.1p2, and are licensed under the 2 point BSD license.

loginrec is written primarily by Andre Lucas, atomicio.c by Theo de Raadt.

strcat() is (c) Todd C. Miller

=====

Import code in keyimport.c is modified from PuTTY's import.c, licensed as follows:

PuTTY is copyright 1997-2003 Simon Tatham.

Portions copyright Robert de Bath, Joris van Rantwijk, Delian Delchev, Andreas Schultz, Jeroen Massar, Wez Furlong, Nicolas Barry, Justin Bradford, and CORE SDI S.A.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Sun Community Source License

SUN COMMUNITY SOURCE LICENSE Version 2.8 (Rev. Date January 17, 2001)

RECITALS

Original Contributor has developed Specifications and Source Code implementations of certain Technology; and

Original Contributor desires to license the Technology to a large community to facilitate research, innovation and product development while maintaining compatibility of such products with the Technology as delivered by Original Contributor; and

Original Contributor desires to license certain Sun Trademarks for the purpose of branding products that are compatible with the relevant Technology delivered by Original Contributor; and

You desire to license the Technology and possibly certain Sun Trademarks from Original Contributor on the terms and conditions specified in this License.

In consideration for the mutual covenants contained herein, You and Original Contributor agree as follows:

AGREEMENT

1. Introduction. The Sun Community Source License and effective attachments ("License") may include five distinct licenses: Research Use, TCK, Internal Deployment Use, Commercial Use and Trademark License. The Research Use license is effective when You execute this License. The TCK and Internal Deployment Use licenses are effective when You execute this License, unless otherwise specified in the TCK and Internal Deployment Use attachments. The Commercial Use and Trademark licenses must be signed by You and Original Contributor in order to become effective. Once effective, these licenses and the associated requirements and responsibilities are cumulative. Capitalized terms used in this License are defined in the Glossary.

2. License Grants.

2.1. Original Contributor Grant. Subject to Your compliance with Sections 3, 8.10 and Attachment A of this License, Original Contributor grants to You a worldwide, royalty-free, non-exclusive license, to the extent of Original Contributor's Intellectual Property Rights covering the Original Code, Upgraded Code and Specifications, to do the following:

a) Research Use License: (i) use, reproduce and modify the Original Code, Upgraded Code and Specifications to create Modifications and Reformatted Specifications for Research Use by You, (ii) publish and display Original Code, Upgraded Code and Specifications with, or as part of Modifications, as permitted under Section 3.1 b) below, (iii) reproduce and distribute copies of Original Code and Upgraded Code to Licensees and students for Research Use by You, (iv) compile, reproduce and distribute Original Code and Upgraded Code in Executable form, and Reformatted Specifications to anyone for Research Use by You.

b) Other than the licenses expressly granted in this License, Original Contributor retains all right, title, and interest in Original Code and Upgraded Code and Specifications.

2.2. Your Grants.

a) To Other Licensees. You hereby grant to each Licensee a license to Your Error Corrections and Shared Modifications, of the same scope and extent as Original Contributor's licenses under Section 2.1 a) above relative to Research Use, Attachment C relative to Internal Deployment Use, and Attachment D relative to Commercial Use.

b) To Original Contributor. You hereby grant to Original Contributor a worldwide, royalty-free, non-exclusive, perpetual and irrevocable license, to the extent of Your Intellectual Property Rights covering Your Error Corrections, Shared Modifications and Reformatted Specifications, to use, reproduce, modify, display and distribute Your Error Corrections, Shared Modifications and Reformatted Specifications, in any form, including the right to sublicense such rights through multiple tiers of distribution.

A

c) Other than the licenses expressly granted in Sections 2.2 a) and b) above, and the restriction set forth in Section 3.1 d)(iv) below, You retain all right, title, and interest in Your Error Corrections, Shared Modifications and Reformatted Specifications.

2.3. Contributor Modifications. You may use, reproduce, modify, display and distribute Contributor Error Corrections, Shared Modifications and Reformatted Specifications, obtained by You under this License, to the same scope and extent as with Original Code, Upgraded Code and Specifications.

2.4. Subcontracting. You may deliver the Source Code of Covered Code to other Licensees having at least a Research Use license, for the sole purpose of furnishing development services to You in connection with Your rights granted in this License. All such Licensees must execute appropriate documents with respect to such work consistent with the terms of this License, and acknowledging their work-made-for-hire status or assigning exclusive right to the work product and associated Intellectual Property Rights to You.

3. Requirements and Responsibilities.

3.1. Research Use License. As a condition of exercising the rights granted under Section 2.1 a) above, You agree to comply with the following:

a) Your Contribution to the Community. All Error Corrections and Shared Modifications which You create or contribute to are automatically subject to the licenses granted under Section 2.2 above. You are encouraged to license all of Your other Modifications under Section 2.2 as Shared Modifications, but are not required to do so. You agree to notify Original Contributor of any errors in the Specification.

b) Source Code Availability. You agree to provide all Your Error Corrections to Original Contributor as soon as reasonably practicable and, in any event, prior to Internal Deployment Use or Commercial Use, if applicable. Original Contributor may, at its discretion, post Source Code for Your Error Corrections and Shared Modifications on the Community Webserver. You may also post Error Corrections and Shared Modifications on a web-server of Your choice; provided, that You must take reasonable precautions to ensure that only Licensees have access to such Error Corrections and Shared Modifications. Such precautions shall include, without limitation, a password protection scheme limited to Licensees and a click-on, download certification of Licensee status required of those attempting to download from the server. An example of an acceptable certification is attached as Attachment A-2.

c) Notices. All Error Corrections and Shared Modifications You create or contribute to must include a file documenting the additions and changes You made and the date of such additions and changes. You must also include the notice set forth in Attachment A-1 in the file header. If it is not possible to put the notice in a particular Source Code file due to its structure, then You must include the notice in a location (such as a relevant directory file), where a recipient would be most likely to look for such a notice.

d) Redistribution.

(i) Source. Covered Code may be distributed in Source Code form only to another Licensee (except for students as provided below). You may not offer or impose any terms on any Covered Code that alter the rights, requirements, or responsibilities of such Licensee. You may distribute Covered Code to students for use in connection with their course work and research projects undertaken at accredited educational institutions. Such students need not be Licensees, but must be given a copy of the notice set forth in Attachment A-3 and such notice must also be included in a file header or prominent location in the Source Code made available to such students.

(ii) Executable. You may distribute Executable version(s) of Covered Code to Licensees and other third parties only for the purpose of evaluation and comment in connection with Research Use by You and under a license of Your choice, but which limits use of such Executable version(s) of Covered Code only to that purpose.

(iii) Modified Class, Interface and Package Naming. In connection with Research Use by You only, You may use Original Contributor's class, interface and package names only to accurately reference or invoke the Source Code files You modify. Original Contributor grants to You a limited license to the extent necessary for such purposes.

(iv) Modifications. You expressly agree that any distribution, in whole or in part, of Modifications developed by You shall only be done pursuant to the term and conditions of this License.

e) Extensions.

(i) Covered Code. You may not include any Source Code of Community Code in any Extensions;

(ii) Publication. No later than the date on which You first distribute such Extension for Commercial Use, You must publish to the industry, on a non-confidential basis and free of all copyright restrictions with respect to reproduction and use, an accurate and current specification for any Extension. In addition, You must make available an appropriate test suite, pursuant to the same rights as the specification, sufficiently detailed to allow any third party reasonably skilled in the technology to produce implementations of the Extension compatible with the specification. Such test suites must be made available as soon as reasonably practicable but, in no event, later than ninety (90) days after Your first Commercial Use of the Extension. You must use reasonable efforts to promptly clarify and correct the specification and the test suite upon written request by Original Contributor.

(iii) Open. You agree to refrain from enforcing any Intellectual Property Rights You may have covering any interface(s) of Your Extension, which would prevent the implementation of such interface(s) by Original Contributor or any Licensee. This obligation does not prevent You from enforcing any Intellectual Property Right You have that would otherwise be infringed by an implementation of Your Extension.

(iv) Class, Interface and Package Naming. You may not add any packages, or any public or protected classes or interfaces with names that originate or might appear to originate from Original Contributor including, without limitation, package or class names which begin with "sun", "java", "javax", "jini", "net.jini", "com.sun" or their equivalents in any subsequent class, interface and/or package naming convention adopted by Original Contributor. It is specifically suggested that You name any new packages using the "Unique Package Naming Convention" as described in "The Java Language Specification" by James Gosling, Bill Joy, and Guy Steele, ISBN 0-201-63451-1, August 1996.

Section 7.7 "Unique Package Names", on page 125 of this specification which states, in part:

"You form a unique package name by first having (or belonging to an organization that has) an Internet domain name, such as "sun.com". You then reverse the name, component by component, to obtain, in this example, "Com.sun", and use this as a prefix for Your package names, using a convention developed within Your organization to further administer package names."

3.2. Additional Requirements and Responsibilities. Any additional requirements and responsibilities relating to the Technology are listed in Attachment F (Additional Requirements and Responsibilities), if applicable, and are hereby incorporated into this Section 3.

4. Versions of the License.

4.1. License Versions. Original Contributor may publish revised versions of the License from time to time. Each version will be given a distinguishing version number.

4.2. Effect. Once a particular version of Covered Code has been provided under a version of the License, You may always continue to use such Covered Code under the terms of that version of the License. You may also choose to use such Covered Code under the terms of any subsequent version of the License. No one other than Original Contributor has the right to promulgate License versions.

5. Disclaimer of Warranty.

5.1. COVERED CODE IS PROVIDED UNDER THIS LICENSE "AS IS," WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED CODE IS FREE OF DEFECTS, MERCHANTABILITY, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. YOU AGREE TO BEAR THE ENTIRE

RISK IN CONNECTION WITH YOUR USE AND DISTRIBUTION OF COVERED CODE UNDER THIS LICENSE. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED CODE IS AUTHORIZED HEREUNDER EXCEPT SUBJECT TO THIS DISCLAIMER.

5.2. You acknowledge that Original Code, Upgraded Code and Specifications are not designed or intended for use in (i) on-line control of aircraft, air traffic, aircraft navigation or aircraft communications; or (ii) in the design, construction, operation or maintenance of any nuclear facility. Original Contributor disclaims any express or implied warranty of fitness for such uses.

6. Termination.

6.1. By You. You may terminate this Research Use license at anytime by providing written notice to Original Contributor.

6.2. By Original Contributor. This License and the rights granted hereunder will terminate: (i) automatically if You fail to comply with the terms of this License and fail to cure such breach within 30 days of receipt of written notice of the breach; (ii) immediately in the event of circumstances specified in Sections 7.1 and 8.4; or (iii) at Original Contributor's discretion upon any action initiated in the first instance by You alleging that use or distribution by Original Contributor or any Licensee, of Original Code, Upgraded Code, Error Corrections or Shared Modifications contributed by You, or Specifications, infringe a patent owned or controlled by You.

6.3. Effect of Termination. Upon termination, You agree to discontinue use and return or destroy all copies of Covered Code in your possession. All sublicenses to the Covered Code which you have properly granted shall survive any termination of this License. Provisions which, by their nature, should remain in effect beyond the termination of this License shall survive including, without limitation, Sections 2.2, 3, 5, 7 and 8.

6.4. Each party waives and releases the other from any claim to compensation or indemnity for permitted or lawful termination of the business relationship established by this License.

7. Liability.

7.1. Infringement. Should any of the Original Code, Upgraded Code, TCK or Specifications ("Materials") become the subject of a claim of infringement, Original Contributor may, at its sole option, (i) attempt to procure the rights necessary for You to continue using the Materials, (ii) modify the Materials so that they are no longer infringing, or (iii) terminate Your right to use the Materials, immediately upon written notice, and refund to You the amount, if any, having then actually been paid by You to Original Contributor for the Original Code, Upgraded Code and TCK, depreciated on a straight line, five year basis.

7.2. LIMITATION OF LIABILITY. TO THE FULL EXTENT ALLOWED BY APPLICABLE LAW, ORIGINAL CONTRIBUTOR'S LIABILITY TO YOU FOR CLAIMS RELATING TO THIS LICENSE, WHETHER FOR BREACH OR IN TORT, SHALL BE LIMITED TO ONE HUNDRED PERCENT (100%) OF THE AMOUNT HAVING THEN ACTUALLY BEEN PAID BY YOU TO ORIGINAL CONTRIBUTOR FOR ALL COPIES LICENSED HEREUNDER OF THE PARTICULAR ITEMS GIVING RISE TO SUCH CLAIM, IF ANY. IN NO EVENT WILL YOU (RELATIVE TO YOUR SHARED MODIFICATIONS OR ERROR CORRECTIONS) OR ORIGINAL CONTRIBUTOR BE LIABLE FOR ANY INDIRECT, PUNITIVE, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH OR ARISING OUT OF THIS LICENSE (INCLUDING, WITHOUT LIMITATION, LOSS OF PROFITS, USE, DATA, OR OTHER ECONOMIC ADVANTAGE), HOWEVER IT ARISES AND ON ANY THEORY OF LIABILITY, WHETHER IN AN ACTION FOR CONTRACT, STRICT LIABILITY OR TORT (INCLUDING NEGLIGENCE) OR OTHERWISE, WHETHER OR NOT YOU OR ORIGINAL CONTRIBUTOR HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE AND NOTWITHSTANDING THE FAILURE OF ESSENTIAL PURPOSE OF ANY REMEDY.

8. Miscellaneous.

8.1. Trademark. You agree to comply with the then current

Sun Trademark & Logo Usage Requirements accessible through the SCSL Webpage. Except as expressly provided in the License, You are granted no right, title or license to, or interest in, any Sun Trademarks. You agree not to (i) challenge Original Contributor's ownership or use of Sun Trademarks; (ii) attempt to register any Sun Trademarks, or any mark or logo substantially similar thereto; or (iii) incorporate any Sun Trademarks into your own trademarks, product names, service marks, company names, or domain names.

8.2. Integration. This License represents the complete agreement concerning the subject matter hereof.

8.3. Assignment. Original Contributor may assign this License, and its rights and obligations hereunder, in its sole discretion. You may assign the Research Use portions of this License to a third party upon prior written notice to Original Contributor (which may be provided via the Community Web-Server). You may not assign the Commercial Use license or TCK license, including by way of merger (regardless of whether You are the surviving entity) or acquisition, without Original Contributor's prior written consent.

8.4. Severability. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. Notwithstanding the foregoing, if You are prohibited by law from fully and specifically complying with Sections 2.2 or 3, this License will immediately terminate and You must immediately discontinue any use of Covered Code.

8.5. Governing Law. This License shall be governed by the laws of the United States and the State of California, as applied to contracts entered into and to be performed in California between California residents. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded.

8.6. Dispute Resolution.

a) Any dispute arising out of or relating to this License shall be finally settled by arbitration as set out herein, except that either party may bring any action, in a court of competent jurisdiction (which jurisdiction shall be exclusive), with respect to any dispute relating to such party's Intellectual Property Rights or with respect to Your compliance with the TCK license. Arbitration shall be administered: (i) by the American Arbitration Association (AAA), (ii) in accordance with the rules of the United Nations Commission on International Trade Law (UNCITRAL) (the "Rules") in effect at the time of arbitration as modified herein; and (iii) the arbitrator will apply the substantive laws of California and United States. Judgment upon the award rendered by the arbitrator may be entered in any court having jurisdiction to enforce such award.

A

b) All arbitration proceedings shall be conducted in English by a single arbitrator selected in accordance with the Rules, who must be fluent in English and be either a retired judge or practicing attorney having at least ten (10) years litigation experience and be reasonably familiar with the technology matters relative to the dispute. Unless otherwise agreed, arbitration venue shall be in London, Tokyo, or San Francisco, whichever is closest to defendant's principal business office. The arbitrator may award monetary damages only and nothing shall preclude either party from seeking provisional or emergency relief from a court of competent jurisdiction. The arbitrator shall have no authority to award damages in excess of those permitted in this License and any such award in excess is void. All awards will be payable in U.S. dollars and may include, for the prevailing party (i) pre-judgment award interest, (ii) reasonable attorneys' fees incurred in connection with the arbitration, and (iii) reasonable costs and expenses incurred in enforcing the award. The arbitrator will order each party to produce identified documents and respond to no more than twenty-five single question interrogatories.

8.7. Construction. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License.

8.8. U.S. Government End Users. The Covered Code is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" and "commercial computer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept. 1995).

Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Code with only those rights set forth herein. You agree to pass this notice to Your licensees.

8.9. Press Announcements. All press announcements relative to the execution of this License must be reviewed and approved by Original Contributor and You prior to release.

8.10. International Use.

a) Export/Import Laws. Covered Code is subject to U.S. export control laws and may be subject to export or import regulations in other countries. Each party agrees to comply strictly with all such laws and regulations and acknowledges their responsibility to obtain such licenses to export, re-export, or import as may be required. You agree to pass these obligations to Your licensees.

b) Intellectual Property Protection. Due to limited intellectual property protection and enforcement in certain countries, You agree not to redistribute the Original Code, Upgraded Code, TCK and Specifications to any country other than the list of restricted countries on the SCSL Webpage.

8.11. Language. This License is in the English language only, which language shall be controlling in all respects, and all versions of this License in any other language shall be for accommodation only and shall not be binding on the parties to this License. All communications and notices made or given pursuant to this License, and all documentation and support to be provided, unless otherwise noted, shall be in the English language.

PLEASE READ THE TERMS OF THIS LICENSE CAREFULLY. BY CLICKING ON THE "ACCEPT" BUTTON BELOW YOU ARE ACCEPTING AND AGREEING TO THE TERMS AND CONDITIONS OF THIS LICENSE WITH SUN MICROSYSTEMS, INC. IF YOU ARE AGREEING TO THIS LICENSE ON BEHALF OF A COMPANY, YOU REPRESENT THAT YOU ARE AUTHORIZED TO BIND THE COMPANY TO SUCH A LICENSE. WHETHER YOU ARE ACTING ON YOUR OWN BEHALF, OR REPRESENTING A COMPANY, YOU MUST BE OF MAJORITY AGE AND BE OTHERWISE COMPETENT TO ENTER INTO CONTRACTS. IF YOU DO NOT MEET THIS CRITERIA OR YOU DO NOT AGREE TO ANY OF THE TERMS AND CONDITIONS OF THIS LICENSE, CLICK ON THE REJECT BUTTON TO EXIT.

GLOSSARY

1. "Commercial Use" means any use (excluding Internal Deployment Use) or distribution, directly or indirectly of Compliant Covered Code by You to any third party, alone or bundled with any other software or hardware, for direct or indirect commercial or strategic gain or advantage, subject to execution of Attachment D by You and Original Contributor.
2. "Community Code" means the Original Code, Upgraded Code, Error Corrections, Shared Modifications, or any combination thereof.
3. "Community Webserver(s)" means the webserver(s) designated by Original Contributor for posting Error Corrections and Shared Modifications.
4. "Compliant Covered Code" means Covered Code that complies with the requirements of the TCK.
5. "Contributor" means each Licensee that creates or contributes to the creation of any Error Correction or Shared Modification.
6. "Covered Code" means the Original Code, Upgraded Code, Modifications, or any combination thereof.
7. "Error Correction" means any change made to Community Code which conforms to the Specification and corrects the adverse effect of a failure of Community Code to perform any function set forth in or required by the Specifications.
8. "Executable" means Covered Code that has been converted to a form other than Source Code.
9. "Extension(s)" means any additional classes or other programming code and/or interfaces developed by or for You which: (i) are designed for use with the Technology; (ii) constitute an API for a library of computing functions or services; and (iii) are disclosed to third party software developers for the purpose of developing software which invokes such additional classes or other programming code and/or interfaces. The foregoing shall not apply to software development by Your subcontractors to be exclusively used by You.
10. "Intellectual Property Rights" means worldwide statutory and common law rights associated solely with (i) patents and patent applications; (ii) works of authorship including copyrights, copyright applications, copyright registrations and "moral rights"; (iii) the protection of trade and industrial secrets and confidential information; and (iv) divisions, continuations, renewals, and re-issuances of the foregoing now existing or acquired in the future.
11. "Internal Deployment Use" means use of Compliant Covered Code (excluding Research Use) within Your business or organization only by Your employees and/or agents, subject to execution of Attachment C by You and Original Contributor, if required.
12. "Licensee" means any party that has entered into and has in effect a version of this License with Original Contributor.
13. "Modification(s)" means (i) any change to Covered Code; (ii) any new file or other representation of computer program statements that contains any portion of Covered Code; and/or (iii) any new Source Code implementing any portion of the Specifications.
14. "Original Code" means the initial Source Code for the Technology as described on the Technology Download Site.
15. "Original Contributor" means Sun Microsystems, Inc., its affiliates and its successors and assigns.

A

16. "Reformatted Specifications" means any revision to the Specifications which translates or reformats the Specifications (as for example in connection with Your documentation) but which does not alter, subset or superset the functional or operational aspects of the Specifications.
17. "Research Use" means use and distribution of Covered Code only for Your research, development, educational or personal and individual use, and expressly excludes Internal Deployment Use and Commercial Use.
18. "SCSL Webpage" means the Sun Community Source license webpage located at <http://sun.com/software/communitysource>, or such other url that Original Contributor may designate from time to time.
19. "Shared Modifications" means Modifications provided by You, at Your option, pursuant to Section 2.2, or received by You from a Contributor pursuant to Section 2.3.
20. "Source Code" means computer program statements written in any high-level, readable form suitable for modification and development.
21. "Specifications" means the specifications for the Technology and other documentation, as designated on the Technology Download Site, as may be revised by Original Contributor from time to time.
22. "Sun Trademarks" means Original Contributor's SUN, JAVA, and JINI trademarks and logos, whether now used or adopted in the future.
23. "Technology" means the technology described in Attachment B, and Upgrades.
24. "Technology Compatibility Kit" or "TCK" means the test programs, procedures and/or other requirements, designated by Original Contributor for use in verifying compliance of Covered Code with the Specifications, in conjunction with the Original Code and Upgraded Code. Original Contributor may, in its sole discretion and from time to time, revise a TCK to correct errors and/or omissions and in connection with Upgrades.
25. "Technology Download Site" means the site(s) designated by Original Contributor for access to the Original Code, Upgraded Code, TCK and Specifications.
26. "Upgrade(s)" means new versions of Technology designated exclusively by Original Contributor as an Upgrade and released by Original Contributor from time to time.
27. "Upgraded Code" means the Source Code for Upgrades, possibly including Modifications made by Contributors.
28. "You(r)" means an individual, or a legal entity acting by and through an individual or individuals, exercising rights either under this License or under a future version of this License issued pursuant to Section 4.1. For legal entities, "You(r)" includes any entity that by majority voting interest controls, is controlled by, or is under common control with You.

ATTACHMENT A REQUIRED NOTICES

ATTACHMENT A-1 REQUIRED IN ALL CASES

"The contents of this file, or the files included with this file, are subject to the current version of Sun Community Source License for [fill in name of applicable Technology] (the "License"); You may not use this file except in compliance with the License. You may obtain a copy of the License at <http://sun.com/software/communitysource>. See the License for the rights, obligations and limitations governing use of the contents of the file.

The Original and Upgraded Code is [fill in name of applicable Technology]. The developer of the Original and Upgraded Code is Sun Microsystems, Inc. Sun Microsystems, Inc. owns the copyrights in the portions it created. All Rights Reserved.

Contributor(s):

Associated Test Suite(s) Location:

ATTACHMENT A-2 SAMPLE LICENSEE CERTIFICATION

"By clicking the 'Agree' button below, You certify that You are a Licensee in good standing under the Sun Community Source License, [fill in name of applicable Technology] ("License") and that Your access, use and distribution of code and information You may obtain at this site is subject to the License."

ATTACHMENT A-3 REQUIRED STUDENT NOTIFICATION

"This software and related documentation has been obtained by your educational institution subject to the Sun Community Source License, [fill in name of applicable Technology]. You have been provided access to the software and related documentation for use only in connection with your course work and research activities as a matriculated student of your educational institution. Any other use is expressly prohibited.

THIS SOFTWARE AND RELATED DOCUMENTATION CONTAINS PROPRIETARY MATERIAL OF SUN MICROSYSTEMS, INC, WHICH ARE PROTECTED BY VARIOUS INTELLECTUAL PROPERTY RIGHTS.

You may not use this file except in compliance with the License. You may obtain a copy of the License on the web at <http://sun.com/software/communitysource>."

ATTACHMENT B

Java (tm) Platform, Standard Edition, Java 2 JDK 1.4.2 Source Technology

Description of "Technology"

Java (tm) Platform, Standard Edition, Java 2 JDK 1.4.2 Source Technology as described on the Technology Download Site.

ATTACHMENT C INTERNAL DEPLOYMENT USE

This Attachment C is only effective for the Technology specified in Attachment B, upon execution of Attachment D (Commercial Use License) including the requirement to pay royalties. In the event of a conflict between the terms of this Attachment C and Attachment D, the terms of Attachment D shall govern.

1. Internal Deployment License Grant. Subject to Your compliance with Section 2 below, and Section 8.10 of the Research Use license; in addition to the Research Use license and the TCK license, Original Contributor grants to You a worldwide, non-exclusive license, to the extent of Original Contributor's Intellectual Property Rights covering the Original Code, Upgraded Code and Specifications, to do the following:

- a) reproduce and distribute internally, Original Code and Upgraded Code as part of Compliant Covered Code, and Specifications, for Internal Deployment Use,

A

b) compile such Original Code and Upgraded Code, as part of Compliant Covered Code, and reproduce and distribute internally the same in Executable form for Internal Deployment Use, and

c) reproduce and distribute internally, Reformatted Specifications for use in connection with Internal Deployment Use.

2. Additional Requirements and Responsibilities. In addition to the requirements and responsibilities described under Section 3.1 of the Research Use license, and as a condition to exercising the rights granted under Section 3 above, You agree to the following additional requirements and responsibilities:

2.1. Compatibility. All Covered Code must be Compliant Covered Code prior to any Internal Deployment Use or Commercial Use, whether originating with You or acquired from a third party. Successful compatibility testing must be completed in accordance with the TCK License. If You make any further Modifications to any Covered Code previously determined to be Compliant Covered Code, you must ensure that it continues to be Compliant Covered Code.

ATTACHMENT D COMMERCIAL USE LICENSE

[Contact Sun Microsystems For Commercial Use Terms and Conditions]

ATTACHMENT E TECHNOLOGY COMPATIBILITY KIT

The following license is effective for the Java (tm) Platform, Standard Edition, Java 2 JDK 1.4.2 Technology Compatibility Kit only upon execution of a separate support agreement between You and Original Contributor (subject to an annual fee) as described on the SCSL Webpage. The applicable Technology Compatibility Kit for the Technology specified in Attachment B may be accessed at the Technology Download Site only upon execution of the support agreement.

1. TCK License.

a) Subject to the restrictions set forth in Section 1.b below and Section 8.10 of the Research Use license, in addition to the Research Use license, Original Contributor grants to You a worldwide, non-exclusive, non-transferable license, to the extent of Original Contributor's Intellectual Property Rights in the TCK (without the right to sublicense), to use the TCK to develop and test Covered Code.

b) TCK Use Restrictions. You are not authorized to create derivative works of the TCK or use the TCK to test any implementation of the Specification that is not Covered Code. You may not publish your test results or make claims of comparative compatibility with respect to other implementations of the Specification. In consideration for the license grant in Section 1.a above you agree not to develop your own tests which are intended to validate conformation with the Specification.

2. Requirements for Determining Compliance.

2.1. Definitions.

a) "Added Value" means code which:

(i) has a principal purpose which is substantially different from that of the stand-alone Technology;

(ii) represents a significant functional and value enhancement to the Technology;

(iii) operates in conjunction with the Technology; and

(iv) is not marketed as a technology which replaces or substitutes for the Technology.

b) "Java Classes" means the specific class libraries associated with each Technology defined in Attachment B.

c) "Java Runtime Interpreter" means the program(s) which implement the Java virtual machine for the Technology as defined in the Specification.

d) "Platform Dependent Part" means those Original Code and Upgraded Code files of the Technology which are not in a share directory or subdirectory thereof.

e) "Shared Part" means those Original Code and Upgraded Code files of the Technology which are identified as "shared" (or words of similar meaning) or which are in any "share" directory or subdirectory thereof, except those files specifically designated by Original Contributor as modifiable.

f) "User's Guide" means the users guide for the TCK which Original Contributor makes available to You to provide direction in how to run the TCK and properly interpret the results, as may be revised by Original Contributor from time to time.

2.2. Development Restrictions. Compliant Covered Code:

a) must include Added Value;

b) must fully comply with the Specifications for the Technology specified in Attachment B;

c) must include the Shared Part, complete and unmodified;

d) may not modify the functional behavior of the Java Runtime Interpreter or the Java Classes;

e) may not modify, subset or superset the interfaces of the Java Runtime Interpreter or the Java Classes;

f) may not subset or superset the Java Classes;

g) may not modify or extend the required public class or public interface declarations whose names begin with "java", "javax", "jini", "net.jini", "sun.hotjava", "COM.sun" or their equivalents in any subsequent naming convention;

h) Profiles. The following provisions apply if You are licensing a Java Platform, Micro Edition Connected Device Configuration, Java Platform, Micro Edition Connected Limited Device Configuration and/or a Profile:

(i) Profiles may not include an implementation of any part of a Profile or use any of the APIs within a Profile, unless You implement the Profile in its entirety in conformance with the applicable compatibility requirements and test suites as developed and licensed by Original Contributor or other authorized party. "Profile" means: (A) for Java Platform, Micro Edition Connected Device Configuration, Foundation Profile, Personal Profile or such other profile as may be developed under or in connection with the Java Community Process or as otherwise authorized by Original Contributor; (B) for Java Platform, Micro Edition Connected Limited Device Configuration, Java Platform, Micro Edition, Mobile Information Device Profile or such other profile as may be developed under or in connection with the Java Community Process or as otherwise authorized by Original Contributor. Notwithstanding the foregoing, nothing herein shall be construed as eliminating or modifying Your obligation to include Added Value as set forth in Section 2.2(a), above; and

A

(ii) Profile(s) must be tightly integrated with, and must be configured to run in conjunction with, an implementation of a Configuration from Original Contributor (or an authorized third party) which meets Original Contributor's compatibility requirements. "Configuration" means, as defined in Original Contributor's compatibility requirements, either (A) Java Platform, Micro Edition Connected Device Configuration; or (B) Java Platform, Micro Edition Connected Limited Device Configuration.

(iii) A Profile as integrated with a Configuration must pass the applicable TCK for the Technology.

2.3. Compatibility Testing. Successful compatibility testing must be completed by You, or at Original

Contributor's option, a third party designated by Original Contributor to conduct such tests, in accordance with the User's Guide. A Technology must pass the applicable TCK for the Technology. You must use the most current version of the applicable TCK available from Original Contributor one hundred twenty (120) days (two hundred forty [240] days in the case of silicon implementations) prior to: (i) Your Internal Deployment Use; and (ii) each release of Compliant Covered Code by You for Commercial Use. In the event that You elect to use a version of Upgraded Code that is newer than that which is required under this Section 2.3, then You agree to pass the version of the TCK that corresponds to such newer version of Upgraded Code.

2.4. Test Results. You agree to provide to Original Contributor or the third party test facility if applicable, Your test results that demonstrate that Covered Code is Compliant Covered Code and that Original Contributor may publish or otherwise distribute such test results.

Upgrading from Version 4 software to Version 5

In this appendix

- AP650..... 763
- AP7131 / AP71xx..... 764
- Controllers 766

AP650

Upgrade Version 4.x to 5.x

A version 5.x controller can upgrade an AP650 running version 4.x to version 5.x using the WISPe upgrade. This capability is enabled using the `legacy-auto-update` command for the controller, either under the device or the profile. The controller adopts the access point using the standard WISPE protocol messages and then downloads the updated image to it, which converts the AP650 to version 5.x version of code.

`Legacy-auto-update` is enabled by default, so updating a controller from 4.x to 5.x will automatically convert associated AP650 units unless this feature is disabled on the wireless controller.

The example below shows how to enable the `legacy-auto-update` command on a RFS4000 controller:

```
rfs4000-22A136# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
rfs4000-22A136(config)#profile rfs4000 default-rfs4000
rfs4000-22A136(config-profile-default-rfs4000)#legacy-auto-update
rfs4000-22A136(config-profile-default-rfs4000)#commit
rfs4000-22A136(config-profile-default-rfs4000)#
```

Downgrade Version 5.x to 4.x

The AP650 can be automatically downgraded to 4.x by connecting it to a controller running the version 4.x. The AP650 tries to discover both 4.x as well as 5.x controllers by default. If it does not find a 5.x controller, but does find a 4.x controller, then it will attempt to adopt to the version 4.x controller. As part of the adoption process the version 4.x controller will download a 4.x image to the AP650.

AP650 Connectivity in Version 5.x

If the AP650 was Layer 2 adopted, or was using DHCP options to adopt to the controller, it will connect to the controller using the same method even after converting to version 5.x.

B

If the AP650 had a static IP address or the controller was specified using a static IP address, that information is currently not maintained during the upgrade to version 5.x. The AP650 will revert to a factory default setting of VLAN1 and DHCP.

AP7131 / AP71xx

Upgrade version 4.x to 5.x

There are two issues to keep in mind when migrating from 4.x to 5.x on the AP71xx:

1. 4.x and 5.x use incompatible image formats. A special migration image is needed to move between the two versions. This migration image is essentially a version 5.x image packaged to look like a 4.x image.
2. Older versions of 4.x image have issues that can leave the CPU clock running at a lower than the normal speed. This causes the version 5.x codebase to stop at its bootos stage. This is recoverable from the serial console by typing in a recovery command and then upgrading again. Newer versions of the 4.x firmware have this fixed starting with the 4.1.1-18R version, which is now the supported minimum image.

Administrators must upgrade their 4.x AP71xx to the minimum supported version of AP firmware (4.1.1-18R). Once the access point is at the minimum supported version the administrator may then upgrade their controllers to 5.x using the migration image. On the 5.x controller administrators need to copy the migration image to `flash:/` and configure that path in the `legacy-auto-update ap71xx cli` command. Once completed when the access point tries to get adopted by the controller, the controller will upgrade the AP to the version 5.x image using the WISPh protocol. After this upgrade the access point will reload and come up with version 5.x firmware and get adopted to the controller following version 5.x mechanisms.

Downgrade 5.x to 4.x

To downgrade an access point running 5.x back to 4.x the reverse migration image from 5.x needs to be used. This image is installed on the AP just as a regular 5.x firmware is installed using ap-upgrade from CLI or UI. Once the access point has been downgraded to 4.x downgrade the controller to 4.x and the access point can be adopted.

All configuration from 5.x are lost when the AP is reverted to 4.x. The original 4.x configuration, if any, will still be present on the access point.

Migration Files

Two migration files released as part of 5.x are:

- 5.x -> 4.x: AP7131-5.1.0.0-xyzR-04010100018R.img
- 4.x -> 5.x: AP7131-5.1.0.0-xyzR.bin

Recovery of a Pre-4.1.1-18R Access Point

If an AP71xx is upgraded using the AP migration image and is running a version of code older than 4.1.1-18R, there is a possibility the AP might get stuck when booting with 5.x in the BootOS. If this occurs, the following error message is seen on the console after bootup:

```
*** cpu not running at correct speed. expected(500Mhz) current> speed(600Mhz) ***
```

If this does occur, a workaround CLI has been added to the AP71xx bootloader:

1. From the AP71xx bootloader type: `achip fix-cpu-speed`
2. Once the access point boots to runtime mode, upgrade again using the `upgrade` command with the latest 5.1 release. This will fix the CPU speed and from this point on, the AP will boot properly on subsequent reboots.

AP7131 / 7131N Connectivity in version 5.x

If the AP7131 was Layer 2 adopted, or was using DHCP options to adopt to the controller, it will connect to the controller using the same method even after converting to version 5.x.

If the AP7131 had a static IP address or the controller was specified using a static IP address, then the configuration restoration process described below ensures there is connectivity.

Configuration Restoration

Some of the configuration items from a 4.X AP7131 are translated and migrated over to the 5.x configuration file after update. These items help ensure connectivity with the controller. The configuration items that are migrated are:

- Hostname
- Port physical configuration (speed, duplex)
- Port L2 configuration (trunking info)
- IP address of controller (translated to 'controller host' in 5.x)
- WAN interface IP addressing
- LAN interface / Subnet1 IP address

If the configuration can not be read properly (bad blocks, any exceptions while reading the NAND etc) then the access point will come up with default 5.x configuration and create a logfile called `legacyapn_<version>.dump.tar.gz` in `flash:/crashinfo` indicating what was migrated and what the errors were for post-analysis.

Version 5.1 Minimum Requirements

The minimum firmware supported for migration to version 5.x is:

- 4.1.1.0-18R.img

B

Controllers

Upgrade Version 4.x to 5.x

Upon upgrading from version 4.x to version 5.x the 5.x controller will save the configuration from 4.x to a new file on flash. The controller's startup-config will point to the version 5.x default startup-config. No configuration is automatically migrated from the 4.x controller. Configuration of the version 5.x settings must be done manually.

The old configuration file from 4.x is renamed to startup-config-wing4 on flash. The password encryption file is also moved to /etc2/encrypt-passwd-wing4.

Migrating Version 4.x Configuration File to 5.2 Configuration File

Brocade provides a tool to convert legacy 4.x configuration files into 5.2 compatible configuration files. This tool is available as a python script or 32-bit windows executable.

To execute the python script use the following syntax:

```
cfgconvert.py <4.x config file> <4.x syntax file> <5.x config file> <options>
```

Parameter	Usage
4.x Config File	The configuration file from the Version 5 device running-config.
4.x Syntax File	The file which contains the 4.x configuration syntax.
5.2 Config File	The Version 5.2 configuration file.
Options	"FULL" or "PARTIAL". FULL means convert the entire configuration. PARTIAL means convert "connectivity" specific configuration items.

To execute the windows batch file run the following:

```
cfgcv <4.x config file>
```

This will produce 5.2 configuration file named <4.x config file>.out

Downgrade Version 5.x to 4.x

Upon downgrading from version 5.x to version 4.x the switch will save the 5.x configuration which is moved to hidden files of the same name (/etc2/.encrypt-passwd-wing5 and /etc2/nvram/.startup-config-wing5). Any previously saved version 4.x configuration, if present, is restored during the downgrade process.

Version 5.2 Minimum Requirements

The following versions of controller and software will be supported for migration to version 5.2:

- 4.x controller source tree: 4.3.1.0-016R, 4.3.2.0-012R, 4.3.3.0-004R
- 5.0 on RFS40xx
- 5.0.3.0-001R, 5.0.1.0-044R, 5.0.2.0-036B on all other controllers