



# Optimizing Aruba WLANs for Roaming Devices

Version 3.3



Solution Guide

## Copyright

© 2009 Aruba Networks, Inc. AirWave®, Aruba Networks®, Aruba Mobility Management System®, Bluescanner, For Wireless That Works®, Mobile Edge Architecture®, People Move. Networks Must Follow®, RFProtect®, The All Wireless Workplace Is Now Open For Business, Green Island, and The Mobile Edge Company® are trademarks of Aruba Networks, Inc. All rights reserved. All other trademarks are the property of their respective owners.

## Open Source Code

Certain Aruba products include Open Source software code developed by third parties, including software code subject to the GNU General Public License (“GPL”), GNU Lesser General Public License (“LGPL”), or other Open Source Licenses. The Open Source code used can be found at this site:

[http://www.arubanetworks.com/open\\_source](http://www.arubanetworks.com/open_source)

## Legal Notice

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

## Warranty

This hardware product is protected by the standard Aruba warranty of one year parts/labor. For more information, refer to the ARUBACARE SERVICE AND SUPPORT TERMS AND CONDITIONS.



[www.arubanetworks.com](http://www.arubanetworks.com)

1344 Crossman Avenue  
Sunnyvale, California 94089

Phone: 408.227.4500  
Fax 408.227.4550

<b>Chapter 1</b>	<b>Introduction</b>	<b>7</b>
	Aruba Reference Architectures	7
	Solution Guide Assumptions and Scope	8
	Design Validation and Testing	9
<b>Chapter 2</b>	<b>Understanding Mobility Use Cases</b>	<b>11</b>
	Device Mobility Categories	11
	Mobile Application Categories	12
	Multi-Purpose Devices	12
	Single-Purpose Devices	13
	Voice Devices	13
<b>Chapter 3</b>	<b>Understanding Design Principles for Roaming Devices</b>	<b>15</b>
	Incremental Design Summary	15
	Design Summary Matrix	16
	Device Configuration	18
	General Device Configuration for All HMD Classes	18
	Configure Optimal HMD Environment Device Settings	18
	Shared or Dedicated SSIDs	18
	Enable 802.11h (Transmit Power Control)	19
	Single-Purpose Device Specific Configuration	19
	Maximize Single-Purpose HMD Battery Life	19
	Voice-Specific Device Configuration	19
	Maximize Handset Battery Life	20
	End-to-End QoS Design	20
	LAN and WAN Performance Baseline	21
	Airtime Optimization	21
	General Airtime Optimization for All HMD Classes	21
	Wireless Load Balancing	22
	Restricting Broadcasts and Multicasts	25
	Limiting “Chatty” Protocols	26
	Voice-Specific Device Configuration	26
	Capacity Planning	27
	Roaming Optimization	28
	General Roaming Optimization for All HMD Classes	28
	Wi-Fi Coverage	29
	VLAN Pooling	29
	Fast-Roaming Technologies (802.11r and OKC)	30
	Single-Purpose Device Specific Configuration	32
	Voice-Specific Device Configuration	32
	Enable ARM with Voice-Aware and Min/Max Output Power	32
	Configure Max Retries, Max Transmit Failures, and	
	Disable Probe Retries	33

	IP Mobility Configuration	33
	General IP Mobility Configuration for All HMD Classes	33
	Choosing Between Layer 2 and Layer 3 Mobility	34
	Layer 2 (VLAN) Mobility	35
	Layer 2 (VLAN) Mobility Design Considerations	35
	Layer 3 (IP) Mobility	36
	Layer 3 (IP) Mobility Design Considerations	36
	IP Multicast Optimization	38
	General Multicast Optimization for All HMD Classes	38
	Multicast Design Considerations	38
	Interference Resistance	40
	General Interference Resistance for All HMD Classes	40
	FHSS and 802.11b/g Co-Existence Design Considerations	40
<b>Chapter 4</b>	<b>Configuring Global Settings for All Roaming Devices</b>	<b>45</b>
	Device Configuration	46
	Configure Optimal HMD Environment Device Settings	46
	Shared or Dedicated SSIDs	46
	Enable 802.11h (Transmit Power Control)	46
	Airtime Optimization	47
	Wireless Load Balancing	47
	Restricting Broadcasts and Multicasts	48
	Limiting “Chatty” Protocols	49
	Roaming Optimization	51
	Wi-Fi Coverage	51
	VLAN Pooling	51
	Fast-Roaming Technologies (OKC)	51
	IP Mobility Configuration	52
	IP Multicast Optimization	52
	Interference Resistance	53
	Physical and RF Design Optimizations	53
	Mode-Aware ARM	53
	Basic Rates	54
	Max-Retries	54
	Noise Immunity	54
<b>Chapter 5</b>	<b>Configuring Incremental Settings for Single-Purpose HMDs</b>	<b>55</b>
	Device Configuration	55
	Roaming Optimization	56
<b>Chapter 6</b>	<b>Configuring Incremental Settings for Voice HMDs</b>	<b>57</b>
	Device Configuration	57
	Maximize Handset Battery Life	57
	End-to-End QoS	58
	Airtime Optimization	58
	Complete a Voice Capacity Plan	58
	Enable Call Admission Control (CAC)	58
	Roaming Optimization	59
	Enable ARM with Voice-Aware and Min/Max Output Power	59
	Configure Max Retries, Max Transmit Failures, and	
	Disable Probe Retries	59
	Max-Retries	59
	Max Transmit Failures	59
	Disable Probe Retries	60

<b>Chapter 7</b>	<b>Configuration Templates for Polycom, Cisco, Vocera, and Ascom</b>	<b>61</b>
	Polycom SpectraLink 8020/8030 Wireless Telephones Configuration Template	61
	Cisco 792x Series Phones Configuration Template	62
	Vocera	63
	Vocera Configuration Template	64
	Ascom i75 Phones Configuration Template	66
<b>Chapter 8</b>	<b>Troubleshooting for Roaming Devices</b>	<b>67</b>
	Scoping the Problem	67
	Mobility Framework	68
	HMD Troubleshooting	68
	Troubleshooting Flow Chart	69
	Symptom #1—Device cannot see any SSIDs	70
	Symptom #2—Device can see some SSIDs but not the one to which it needs to connect	70
	Symptom #3—Device successfully authenticates but cannot communicate	72
	Symptom #4—Device successfully authenticates and can communicate but is experiencing dropped connections and/or poor performance	73
	Before you contact Aruba Support	74
<b>Appendix A</b>	<b>Device Interoperability Matrix</b>	<b>77</b>
	Commonly Deployed Single-Purpose HMDs	77
	Commonly Deployed Voice HMDs	80
<b>Appendix B</b>	<b>Examples of Highly Mobile Devices</b>	<b>81</b>
<b>Appendix C</b>	<b>Mobility Performance Test Results</b>	<b>85</b>
	Network Design Scaling Considerations	85
	Controller Scaling Considerations	86
	Data Device Roaming Considerations	86
	Voice Device Roaming Considerations	86
<b>Appendix D</b>	<b>Aruba Contact Information</b>	<b>89</b>
	Contacting Aruba Networks	89



This Solution Guide describes best practices for implementing an Aruba 802.11 wireless network that supports thousands of highly mobile devices (HMDs) such as Wi-Fi phones, handheld scanning terminals, voice badges, and computers mounted to vehicles. It describes the design principles particular to keeping devices that are in constant motion connected to the network as well as best practices for configuring Aruba Networks controllers and the mobile devices. The comprehensive guide addresses six areas of network planning to ensure a high quality of service for roaming data and voice sessions: device configuration, airtime optimization, roaming optimization, IP mobility configuration, IP multicast configuration, and interference resistance. A detailed troubleshooting section covers common issues that arise with these types of WLANs.

When a requirement exists for a wireless network to support many roaming devices, it has a significant impact on six areas of network planning:

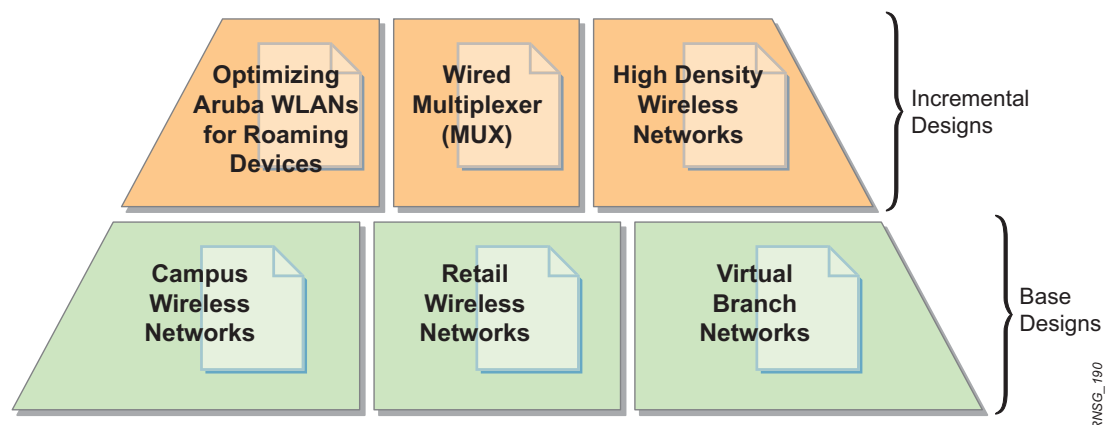
- Device Configuration
- Airtime Optimization
- Roaming Optimization
- IP Mobility Configuration
- IP Multicast Configuration
- Interference Resistance

This guide is divided into major sections on design, configuration, and troubleshooting for roaming devices. It begins by explaining the design principles involved in each of these six areas, as well as specific Aruba features of interest to the wireless architect. In the configuration section, we provide practical guidelines in all six areas on how to program Aruba Mobility Controllers as well as the mobile devices themselves. Finally, there is a detailed troubleshooting procedure for common issues encountered by roaming devices.

## Aruba Reference Architectures

Aruba publishes two types of validated reference architectures: Base Designs and Incremental Designs. [Figure 1](#) is a roadmap to the Aruba Validated Reference Design (VRD) library that illustrates the relationship between these two types of documents.

**Figure 1** Aruba Validated Reference Design Library



The VRD presented in this document is an incremental design. It provides a best practice architecture for a large-scale deployment with thousands of HMDs. It is intended to layer on top of one of the other Aruba base reference network designs (such as *Campus Wireless Networks VRD* or *Retail Wireless Networks VRD*) so that the performance of the HMDs is optimized and that the network achieves a high level of performance.

## Solution Guide Assumptions and Scope

This guide is based on Aruba Operating System version 3.3.2.12. It makes assumptions about the knowledge level of the engineer, the existing architecture and configuration of the Aruba WLAN, and the specific types of mobility that will be supported. Please review [Table 1](#) and verify that your environment conforms to these assumptions.

**Table 1** *Solution Guide Assumptions*

Category	Assumption
Prerequisites — Technical Knowledge	<ul style="list-style-type: none"> <li>Comfort with controller-based WLAN architectures that employ thin APs</li> <li>Thorough understanding of Aruba controller design, master/local architectures, controller and AP redundancy</li> <li>Thorough understanding of profile-based configuration in ArubaOS 3.x, and general familiarity with the profile hierarchy</li> <li>General understanding of virtual AP concepts and configuration</li> </ul>
Prerequisites — Base Design	<ul style="list-style-type: none"> <li>The Aruba base network complies with one of the Aruba published VRDs (Campus, Retail, or Remote Networks) using one or more master/local clusters</li> <li>The Aruba WLAN base network has been defined and/or implemented: <ul style="list-style-type: none"> <li>The number of controllers has been defined</li> <li>The number of APs has been defined</li> <li>The redundancy models used for master controllers, local controllers, and APs have been defined</li> <li>Service Set Identifiers (SSIDs) have been chosen</li> <li>Encryption types and Authentication, Authorization, and Accounting (AAA) integration decisions are made</li> </ul> </li> <li>Complete control over the RF airspace; freedom to choose any combination of channels and power levels that are legal within the regulatory domain</li> </ul>
Included Use Cases	<ol style="list-style-type: none"> <li>Client population: <ul style="list-style-type: none"> <li>Between 100 - 10,000 wireless HMD clients per facility</li> <li>HMDs are over 25% of the wireless client population</li> <li>Moving up to a maximum of 15 mph</li> </ul> </li> <li>Authentication and security: <ul style="list-style-type: none"> <li>Encryption: WPA or WPA2</li> <li>Authentication: PSK or 802.1X</li> </ul> </li> <li>Session types: <ul style="list-style-type: none"> <li>Voice codec data streams and control sessions</li> <li>Character-based and thin-client data applications communicating with onsite or remote application servers</li> <li>Multicast-based machine-to-machine telemetry applications</li> <li>Video applications streaming to mobile devices</li> </ul> </li> </ol>
Excluded Use Cases	<p>This solution guide does not consider the following:</p> <ul style="list-style-type: none"> <li>Ultra dense AP deployments (less than 60 foot spacing between APs)</li> <li>Active RFID or Real-Time Location System (RTLS) tags</li> <li>Guest Access via captive portal authentication</li> <li>Fixed-Mobile Convergence (FMC) or Unified Mobile Access (UMA) handsets</li> <li>PBX/Call Manager Integration</li> </ul>



## Design Validation and Testing

An Aruba VRD packages network designs, deployment methodologies, configuration procedures, and detailed descriptions of product functionality, serving as a reference model for common customer deployment scenarios. Each Aruba VRD is based on best practices derived from large-scale customer deployments. VRD designs are then constructed in a lab environment and thoroughly tested by Aruba engineers. By using these proven designs, Aruba customers can deploy solutions rapidly, with the assurance that they will perform and scale as expected.

Test cases for this VRD were executed against the physical architecture recommended in this guide using a mix of client devices and interconnect methods. ArubaOS release 3.3.2.12 was used to conduct these tests. Please see [Appendix C, “Mobility Performance Test Results”](#) for further details on test case results.



This chapter provides a framework to organize the many types of mobile devices and applications that a wireless network may support. Such a framework is needed to understand the specific issues and challenges presented by varying combinations of mobility requirements. In the remainder of this guide, we use this framework to discuss how you can optimize controllers and APs to create a WLAN that effectively handles HMDs.




## Device Mobility Categories

A WLAN must be flexible enough to accommodate wireless devices that have varying degrees of mobility. Devices can be grouped into three mobility categories based on their usage characteristics and roaming frequency:

- Stationary devices (SDs)
- Somewhat mobile devices (SMDs)
- Highly mobile devices (HMDs) or roaming devices

A brief summary of each device category is given in [Table 2](#). For detailed information about and photos of HMDs, see [Appendix B, “Examples of Highly Mobile Devices”](#).

**Table 2** Examples of Different Mobility Levels for Wireless Devices

Stationary Devices (SDs)	Somewhat Mobile Devices (SMDs)	Highly Mobile Devices (HMDs)
 <p><b>Example: Wireless Scale</b></p> <p>If a wireless network uses only stationary devices, such as wireless printers, scales, or fixed Point of Sale (POS) terminals, the network planning process becomes relatively simple.</p> <p>SDs typically have the following attributes:</p> <ul style="list-style-type: none"> <li>• The device does not roam at all</li> <li>• Bidirectional signal strength is relatively constant</li> <li>• If there is RF interference, it can be identified and solved during initial deployment</li> <li>• Device density is usually sparse</li> </ul>	 <p><b>Example: Patient Monitor</b></p> <p>In some settings, wireless devices move infrequently, or they may move regularly, but are only used while stationary. Examples include laptops used in auditoriums or conference rooms, or patient monitors in hospitals that are used at a bedside.</p> <p>Challenges presented by SMDs include:</p> <ul style="list-style-type: none"> <li>• The device roams, but infrequently</li> <li>• IP multicast may be required by devices such as patient monitors</li> <li>• May require Layer 3 (L3) mobility, depending on the topology of the network</li> <li>• Occasional RF interference</li> <li>• Large numbers of laptops or other devices concentrated in one room may exceed per-AP client service limits</li> </ul>	 <p><b>Example: Wi-Fi Phone</b></p> <p>HMDs are the most difficult WLAN devices to plan for and to implement. HMDs are the subject of this guide. HMDs such as handheld scanners or voice handsets present the following challenges to a WLAN:</p> <ul style="list-style-type: none"> <li>• Continuous roaming events</li> <li>• Device is in use while roaming</li> <li>• APs must continually load balance clients</li> <li>• APs must provide consistent performance across a dynamic range of received signal strengths</li> <li>• Users and applications expect roaming transitions to be undetectable</li> <li>• Devices are more likely to encounter RF interference</li> </ul>

## Mobile Application Categories

Wireless devices can also be divided into categories based on the number and type of applications they support. Aruba defines three broad classes:

- **Multi-purpose:** A laptop or other portable data device that has several applications running within a general-purpose windowing operating system, typically with a high-performance CPU and a high power radio.
- **Single-purpose:** A purpose-built device that runs a single primary application. This may be a thin-client application, but in most cases the browser cannot be used for general Internet access or email. Such a device typically uses a slower CPU and a lower-power radio to conserve battery life.
- **Voice:** A voice handset is a special case of a single-purpose device. Voice devices operate in single or dual mode. Single mode phones operate exclusively over either a private wireless LAN or a public cellular network, while dual-mode handsets can operate in Wi-Fi or Cellular mode.

These application categories can be combined in matrix fashion with the mobility categories mentioned above. [Table 3](#) presents such a matrix, with common examples from a variety of customer environments. The balance of this section describes the three application categories in more detail.

**Table 3** *Device and Application Category Matrix*

		Device Mobility Categories		
		Stationary Devices (SDs)	Somewhat Mobile Devices (SMDs)	Highly Mobile Devices (HMDs)
Mobile Application Categories	Multi-Purpose Device	PC	Laptop	<ul style="list-style-type: none"> <li>• Workstation on Wheels (WOW)</li> <li>• Tablet PC</li> <li>• Ultra-Mobile PC (UMPC)</li> <li>• Smartphone</li> </ul>
	Single-Purpose Device	<ul style="list-style-type: none"> <li>• Wireless copier or fax</li> <li>• Wireless projector</li> <li>• Wireless scale</li> </ul>	<ul style="list-style-type: none"> <li>• Patient monitors</li> <li>• Infusion pumps</li> <li>• Blood pressure machines</li> <li>• Heart monitors</li> <li>• Barcode scanners</li> </ul>	<ul style="list-style-type: none"> <li>• Handheld scanning terminal</li> <li>• Mobile printer</li> <li>• Vehicle-mounted data terminal</li> <li>• Industrial robot</li> <li>• 802.11 RTLS Tag</li> </ul>
	Voice Device	<ul style="list-style-type: none"> <li>• IP desk phone</li> <li>• IP video camera</li> </ul>	n/a	<ul style="list-style-type: none"> <li>• 802.11 voice handset</li> <li>• 802.11 voice badge</li> <li>• Unified Mobile Access (UMA) handset</li> <li>• Fixed Mobile Convergence (FMC) handset</li> </ul>

### Multi-Purpose Devices

Multi-purpose devices are capable of running multiple applications simultaneously, using a window-based operating system to provide core services and networking support. The most versatile HMD is the common laptop. However, new hardware form factors such as the tablet PC, the Ultra-Mobile PC (UMPC) and others have enabled multi-purpose devices to be deployed in novel ways. For example, in the healthcare industry, devices such as workstations on wheels (WOWs) and tablet PCs allow healthcare professionals to keep moving while using these devices to check and monitor patients, write prescriptions, and view lab results.

The roaming behavior of multi-purpose devices can vary widely between devices. This is a function not so much of the device itself, but of the complexity of the various elements that must be integrated together to create a complete multi-purpose platform:

- The wireless Network Interface Card (NIC) chipset and radio
- The NIC driver vendor and revision
- The NIC antenna (single, dual, diversity)
- The NIC wireless support (802.11a/b/g/n)
- The supplicant being used to configure the SSID name and encryption type
- The device operating system and network protocol stack

Some of the more popular operating systems for laptop form-factor devices are Windows XP, Windows Vista, and MacOS, but they don't all support wireless in the same fashion. For example, MacOS does not support opportunistic key caching (OKC). Therefore, the SSID that the MacOS laptops will associate to should either have OKC enabled with validate pmkid or have OKC disabled in the 802.1X profile to provide better roaming mobility. There are a variety of WLAN supplicants that could be used with laptops running Microsoft Windows OS, depending on the laptop manufacturer or wireless NIC that is installed. Each of these supplicants may affect roaming behavior due to their feature support. For example, the Dell supplicant with Broadcom chipset does not support OKC. However, the Microsoft Windows Zero Config supplicant does support OKC.

More recently, a new form factor of multi-purpose HMD has been introduced—the “smartphone.” These devices, such as the Apple iPhone or Research In Motion BlackBerry, allow you to access files or otherwise run applications both in an 802.11 environment and over a public 3G carrier network. Aruba classifies these devices as multi-purpose HMDs because their traffic profiles more closely resemble window-based devices than voice or single-mode devices.

## Single-Purpose Devices

A single-purpose or “purpose-built” mobile WLAN client is designed to perform a specific function, which may be required anywhere in a customer facility. These clients include an 802.11 radio and assume pervasive availability of 802.11 wireless service to do their job.

In addition to the intended function (application) performed by the purpose-built client, it may also be desirable for the wireless network to provide certain information about that client, such as its location. Such “auxiliary” or “support application” information provides incremental value to the primary function. If such auxiliary applications are desired, additional network design choices may be required to enable this support.

Some examples of purpose-built devices are:

- **Healthcare:** Infusion pumps, blood pressure monitors, barcode scanners
- **Commercial printing:** Robotic paper roll movers
- **Retail:** Robotic warehouse picking devices

## Voice Devices

Single-mode and dual-mode voice handsets are a special case of a single-purpose device. Due to the many unique network architecture and configuration requirements for voice service, these are best treated separately. Later in this guide, we explain how to decide whether voice devices need a special SSID or whether they can be combined onto a single SSID with other devices.

A single-mode voice handset has one radio, typically either 802.11a/b/g or a 2.5G/3G CDMA or GSM radio for use with a paid cellular carrier subscription. Dual-mode phones integrate two radios, so that they can roam between cellular networks and voice-over-wireless LANs. Both single-mode and dual-

mode handsets may include additional radios such as Bluetooth for personal area network peripherals such as headsets.



---

The scope of this solution guide focuses on single-mode voice devices.

---

Wireless single-mode 802.11 phones from companies such as Avaya, Polycom, Ascom, Vocera, and Cisco are designed for high-use enterprise environments such as healthcare, in which employee mobility, responsiveness, and productivity are essential.

The handsets allow for decreased use of paging and more expensive cellular connections. Some handset models are more durable than cellular devices, and may also be designed for easy sterilization in clinical environments.

This chapter details the general design principles that a wireless architect must apply to the base design when planning a WLAN to work properly with large numbers of roaming devices.

## Incremental Design Summary

The purpose of this document is to describe the incremental changes that an architect must make to one of the base Aruba reference designs so that it is optimized for an HMD environment. These changes fall into six areas and can be summarized as follows:

- **Device Configuration**

Mobile device default values are not always the best choice for roaming in dense WLANs. Some device changes require that corresponding changes be made on the WLAN infrastructure. Common examples include basic rate support, DTIM settings, and roaming threshold triggers.

- **Airtime Optimization**

Roaming devices are acutely sensitive to RF congestion and inefficiencies. Device performance can be greatly improved by using proper load balancing across APs and channels, leveraging the airtime fairness feature of Aruba's Adaptive Radio Management (ARM) technology, employing client load balancing across the RF spectrum (thus limiting "chatty" protocols such as mDNS and SSDP via a firewall ACL), employing multicast rate optimization, and other techniques.

- **Roaming Optimization**

Roam decisions can be influenced by the infrastructure by optimizing data rates, output power, retry thresholds, and by using the Handoff Assist feature. Roam event durations can be reduced with opportunistic key caching (OKC), 802.1X FastConnect, and proper VLAN design.

- **IP Mobility Configuration**

ArubaOS provides seamless wireless connectivity as users roam between access points using the Mobile IP feature. While rarely a factor in a base design, good IP Mobility design is critical to HMD environments. The selection of layer-2 (VLAN Mobility) or layer-3 (Mobile IP) roaming requires careful planning. Designing mobility domains to minimize inter-controller roaming events can increase overall roaming times.

- **IP Multicast Optimization**

Reducing multicast traffic over the air and on the wire is vital to channel efficiency. Multicast traffic is steadily increasing on WLANs, and is normally transmitted at the lowest supported data rate. Aruba offers a "smart" multicast rate optimization to use higher rates where possible. The Aruba IGMP snooping feature eliminates wire-side multicast frame forwarding to APs that have no multicast users.

- **Interference Resistance**

Because HMDs move more than any other WLAN client type, they are the most likely to encounter and be impacted by adverse RF conditions. Customers who depend on HMDs and who have environments that are known to have serious RF interference must implement additional measures to protect performance. Such measures include output power controls, using lower data rates to avoid OFDM modulations, and enabling interference resistance features.

Device

Airtime

Roaming

IP Mobility

IP Multicast

Interference

The design principles include using Aruba features that provide the wireless architect with additional control over the client, or the ability to mitigate the types of environmental problems that are commonly encountered.

## Design Summary Matrix

The discussion that follows touches on many different device attributes, ArubaOS features, ArubaOS configuration profiles, and even installation decisions. To help you navigate the rest of this document, [Table 4](#) distills all of the main design changes for roaming devices.

Use [Table 4](#) on the following page to keep the design detail in perspective and keep track of which parts of the system are being discussed at various points in the text. The table maps key elements of the base network design to the corresponding mobility optimizations.

Here's how to read the table. There is one row for each of the design areas. The columns contain information as follows:

- **Physical and RF Design Optimizations:** The first cell in each row contains any physical changes to the environment required for that area *that are external to the controller*.
- **Global CLI Config Optimizations:** Startup configuration file changes that are global (that is, external to Aruba profiles) are listed here.
- **Virtual AP Profile:** Changes suggested in the ArubaOS Virtual AP Profile are listed here.
- **SSID Profile:** Changes to the ArubaOS SSID Profile are shown here.
- **AAA Profile:** Recommended changes in the ArubaOS AAA Profile are shown next.
- **RF Radio Profile:** Changes to the ArubaOS RF Radio Profile are listed here.
- **RF ARM Profile:** Changes to the ArubaOS RF ARM Profile are shown here.

Certain configuration optimizations are repeated on more than one row, allowing the wireless architect to choose which specific design areas to apply to a given installation.

In the Configuration section, detailed procedures are provided to address each of the recommended changes. CLI snippets are also included in many cases.



**Table 4** Incremental Design Summary Matrix

	Physical and RF Design Optimizations	Controller Optimizations				
		Global CLI Config Optimizations	AP Group Optimizations for SSIDs that Support HMDs (By Profile)			
			Virtual AP Profile	SSID Profile	AAA Profile	RF Radio Profile
Device	<ul style="list-style-type: none"> <li>Configure optimal HMD environment device settings</li> <li>Enable &amp; verify end-to-end QoS on LAN</li> <li>Enable WMM on client</li> <li>Verify LAN throughput</li> </ul>	<ul style="list-style-type: none"> <li>Enable QoS on controller</li> </ul> <b>QoS</b>	<ul style="list-style-type: none"> <li>Non-default VAP settings</li> </ul>	<b>Shared or Dedicated SSIDs</b> <ul style="list-style-type: none"> <li>Set DTIM to manufacturer preferred minimum</li> <li>Non-default SSID settings</li> <li>Enable U-APSD or Aruba Battery Boost</li> </ul>	<ul style="list-style-type: none"> <li>Non-default AAA settings</li> </ul>	<ul style="list-style-type: none"> <li>Enable 802.11h</li> </ul>
Airtime	<ul style="list-style-type: none"> <li>Purchase 5 GHz capable devices</li> </ul>	<ul style="list-style-type: none"> <li>Enable Airtime Fairness</li> </ul>	<ul style="list-style-type: none"> <li>Enable Band Steering</li> <li>Enable Spectrum Load Balancing</li> </ul>	<b>Wireless Load Balancing</b>		<ul style="list-style-type: none"> <li>Enable Mode-Aware ARM</li> </ul>
Roaming	<ul style="list-style-type: none"> <li>Limit "Chatty" Protocols</li> </ul>	<ul style="list-style-type: none"> <li>Configure Eth ACL to disable IPv6</li> <li>Set user role ACL to deny chatty protocols</li> </ul>	<ul style="list-style-type: none"> <li>Enable Proxy Arp</li> <li>Configure drop multicast OR enable multicast rate optimization</li> </ul>	<b>Reduce Broadcast &amp; Multicast</b>		
IP Mobility	<ul style="list-style-type: none"> <li>Complete per-AP voice capacity plan</li> <li>Ensure uniform minimum signal coverage</li> <li>Purchase OKC-compatible devices</li> </ul>	<ul style="list-style-type: none"> <li>Enable CAC</li> </ul> <b>Optimize Voice</b>	<ul style="list-style-type: none"> <li>Use VLAN pooling in large networks.</li> </ul>	<ul style="list-style-type: none"> <li>Remove high data rates</li> </ul>	<b>Facilitate Roaming</b>	
IP Multicast		If using mobile-ip then create Home Agent Table entries for each subnet where seamless roaming is desired.	Enable mobile-ip OR vlan mobility	<ul style="list-style-type: none"> <li>Set Max Retries</li> <li>Set Max Retry Failures</li> </ul>	<ul style="list-style-type: none"> <li>Disable OKC unless device supports it</li> <li>Enable validate pmkid</li> </ul>	<ul style="list-style-type: none"> <li>Increase Min_TX_Power to 18</li> </ul>
Interference	<ul style="list-style-type: none"> <li>Ensure SNR exceeds manufacturer minimum</li> <li>Locate FHSS APs &gt;20 ft from 802.11 APs</li> <li>Phase out interfering devices</li> <li>Purchase 5 GHz HMDs</li> </ul>	<ul style="list-style-type: none"> <li>Increase Noise Immunity to 3, 4 or 5 (requires ArubaOS 3.3.2.11 or later)</li> </ul>	<b>Multicast Enablement</b>		Enable multicast optimization.	<b>CCI &amp; ACI Interference</b> <ul style="list-style-type: none"> <li>Enable ARM for channel &amp; power management</li> <li>Enable Mode-Aware ARM</li> <li>Increase Min_TX_Power to 18</li> </ul>
				<ul style="list-style-type: none"> <li>Use 802.11b only rates:               <ul style="list-style-type: none"> <li>Basic rates of 1 &amp; 2 only</li> <li>Supported rates of 1,2,5.5 &amp; 11 only</li> </ul> </li> <li>Configure DTIM = XX</li> <li>Set Max Retries to 15</li> </ul>	<b>FHSS and Fixed Frequency Interference</b>	

**Table Legend:**

- Orange Bubble = All roaming devices
- Blue Bubble = Voice-only devices
- Green Bubble = Single-mode & voice devices

## Device Configuration

We begin with a review of design principles that apply to all HMD classes (multi-purpose, single-purpose and voice). This review is followed by specific design rules for voice device configuration.

### General Device Configuration for All HMD Classes

This section considers optimal per-HMD default value modifications, when a dedicated SSID may be called for, and the use of 802.11h. These potential configuration changes apply to every type of roaming device.

#### Configure Optimal HMD Environment Device Settings

Every HMD design should begin by identifying sub-optimal default values on the most prevalent roaming devices. While default values are often more than adequate for Stationary Devices and Somewhat Mobile Devices, it is Aruba's experience that most HMDs can benefit from changes to the defaults. This is particularly true for mobile devices produced prior to the widespread adoption of thin APs and dense deployment strategies in the 2005 timeframe.

It is not uncommon for experienced engineers to have different views about which settings produce benefit and what the nominal values may be for those settings.

Aruba recommends using a five-step process for each widely-deployed HMD:

1. Update all devices to the latest available firmware tested by Aruba, or check with the device manufacturer for a recommended version. See [Appendix A, "Device Interoperability Matrix" on page 77](#) for a list of device firmware levels that have undergone interoperability testing at Aruba.
2. Make a list of all configurable settings and default values.
3. Contact the device vendor systems engineer and request a list of the latest known "best practice" settings for your specific deployment scenario.
4. Make the same request of your Aruba systems engineer.
5. Conduct a pilot test to experiment with these values, and with others that seem relevant but may not be known to the systems engineers.

Remember that some HMDs expose different configuration values to different tools. For example, voice handsets typically have a subset of values that can be configured directly on the phone, while a separate provisioning tool provides much deeper control over the device.

*Configuration Areas Affected:*

- Physical and RF Design

#### Shared or Dedicated SSIDs

The next device design decision is whether or not to use a dedicated SSID (and a separate Virtual AP) for one or more of the HMD devices. This choice should be based more on the device's RF and 802.11 capabilities than on security. The dedicated firewall integrated into the controller allows the administrator to isolate the SSID used for connectivity from the security and QoS policies, which are based on the user profile and traffic type.

The wireless architect should always seek to use shared SSIDs unless there is a specific reason to do otherwise. Every defined SSID consumes system resources for policy application, LAN bandwidth due to additional tunnels, and spectrum for beacons and other management overhead. In many cases, the cost of additional SSIDs outweighs the benefits. However, for certain roaming devices, a dedicated SSID can result in significant throughput or battery performance advantages.

A dedicated SSID should be used for HMDs if one of the following applies:

- The majority of devices that will associate to the SSID have manufacturer suggested Delivery Traffic Indication Message (DTIM) settings greater than the default settings. The battery save settings (Power Save and DTIM settings) on some devices like voice and single-purpose HMDs can be optimized to larger DTIM values to improve battery life without adversely affecting the device operations. Changes in the DTIM value affect data client performance, so if this SSID parameter needs to be modified outside of the default settings, then a different SSID profile should be used for that group of HMDs.
- Any non-default 802.11 settings need to be configured to optimize HMD connectivity performance.
- Any non-default AAA profile settings that need to be configured for some HMDs.
- Any non-default Virtual AP profile settings that need to be configured for some HMDs.
- The encryption and authentication levels supported by some HMDs do not match the encryption and authentication mechanisms enforced on multi-purpose HMDs.
- The encryption and authentication methods supported by the HMDs match the security enforced on multi-purpose HMDs, but these settings adversely affect HMD roaming due to possible legacy driver behavior or processing power, which would require different 802.1X profile settings to resolve certain key or timing behavior.
- The HMD device infrastructure, such as voice HMD, demands a dedicated VLAN because it does not support L3 connectivity back to certain application servers (like voice call servers).

If none of these criteria match, it should be possible to use the same SSIDs and the same encryption and authentication methods for all HMDs. Different levels of QoS can be enforced based on the traffic type without requiring a separate SSID.

*Configuration Areas Affected:*

- Virtual AP Profile
- SSID Profile

### Enable 802.11h (Transmit Power Control)

To help provide better connectivity for HMD clients, enable 802.11h in both the 802.11a and 802.11g RF profiles. This setting causes the AP to add the 802.11d country and power constraint information element in the 802.11 header of Beacons and Probe Responses.

*Configuration Areas Affected:*

- RF Radio Profile

## Single-Purpose Device Specific Configuration

This section considers battery life extension for single-purpose devices.

### Maximize Single-Purpose HMD Battery Life

Good battery life is an important usability consideration for any single-purpose HMD, just as it is for voice devices. Single-purpose devices must carefully manage limited battery capacity to maximize the time between recharges.

- Determine the manufacturer's recommended DTIM value for your particular deployment scenario. This value must be implemented on both the controller as well as the device (which may default to another value).

## Voice-Specific Device Configuration

This section considers battery life extension, enabling of QoS from end to end, and creating a LAN performance baseline prior to deploying voice.

## Maximize Handset Battery Life

Good battery life is an important usability consideration for VoFi installations. While much battery-saving technology depends on the handset, the following functions of the infrastructure can assist in extending battery life:

- Determine desired battery life. This function depends on applications. For retail and hospital environments, where people pick up their phone in the morning or at the beginning of their shift and then return it to a charging cradle at the end of the day, a battery life of 12 hours may be adequate. This number is easily achievable today.
- UAPSD is part of the 802.11e specification that the Wi-Fi Alliance is certifying as Wireless MultiMedia Power Save (WMM-PS). This feature offers considerable improvement in battery life over former protocols, and while few handsets support UAPSD today, any new Wi-Fi infrastructure must be UAPSD/WMM-PS capable.
- ARP proxy. The key to extending battery life, particularly idle (not on-call) battery life, is to reduce the LAN traffic the handset sees. The most significant type of traffic is Address Resolution Protocol (ARP), and good WLANs now proxy to reduce this traffic to voice clients.
- Traffic filtering. Some vendors offer features where extraneous traffic is not delivered to voice clients, thus saving the battery drain entailed in receiving and ACK'ing such traffic. Many types of multicast fall into this category.

*Configuration Areas Affected:*

- SSID Profile

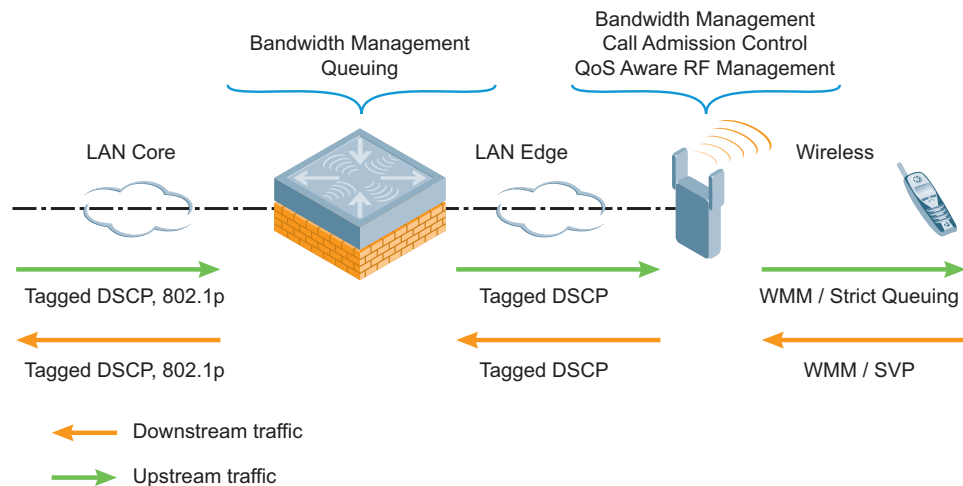
## End-to-End QoS Design

Different applications have different QoS requirements. The International Telecommunication Union (ITU) provides recommendations for good voice quality:

- Round trip delay for voice traffic and call setup control traffic should be less than 100 ms.
- Jitter is the variation in time of the arriving packets due to delays introduced in the network. Each handset vendor has a different tolerance level to jitter, depending on its implementation of jitter buffers. A jitter of more than 10 ms is considered unacceptable.
- Loss of packets often results in dropped audio. Packet loss should be less than 5%.

QoS enforcement is not the responsibility of any one network component. All of the devices in the network should be able to recognize the relative priority of the traffic and prioritize locally accordingly. QoS enforcement also requires client participation, especially over the air.

**Figure 2** End-to-End QoS



Ref\_117

## Wired QoS Recommendations

- Use 802.1p tags and DSCP tags to prioritize the traffic on the wire. Commonly used values for voice are ToS 6 or 7 and dot1p 6.
- If supported, priority queuing for voice traffic should be enabled on the routers and switches.
- All the routers and switches in the network path between APs and the controller should be configured to recognize the tagging on the traffic and prioritize the traffic accordingly.
- If network devices rewrite the tags on the packet headers, make sure that all network devices use the same tags for a given traffic type.

## Wireless QoS Recommendations

- If the client supports WMM, enable WMM on the Aruba system and the handset. Verify that the client tags the voice (and data traffic) appropriately.
- WMM queues map to different DSCP and ToS tags. Unless otherwise recommended by the handset manufacturer, use the default WMM mappings.
- Make sure that the traffic prioritization is such that voice receives the highest priority, followed by video. Data should receive the lowest priority. The priority levels for each of the applications are set according to the delay, retry, jitter, and loss tolerance of the application.
- Make sure that the protocols for voice data traffic (for example, RTP) and control traffic (for example, SIP) are prioritized. Control traffic is used for call setup and the voice data traffic needs to be prioritized to provide good call quality.
- In the absence of WMM support on the handset, make sure that voice uses the high-priority queue and all the other applications use the low-priority queue.

*Configuration Areas Affected:*

- Physical and RF Design

## LAN and WAN Performance Baseline

Because voice is an end-to-end application, it is essential to verify that the LAN (and WAN if traffic is being backhauled between sites) is ready for voice traffic.

If the Enterprise already has VoIP applications on the LAN, there should be minimal extra work required to extend the applications to Voice over Wi-Fi (VoFi).

Start with a quick check of capacity. Assume each voice call takes 100 Kbps each way over the LAN (200 Kbps for both directions). This assumes G.711 and a 64 Kbps CODEC. There is some bandwidth reduction if compressed voice is used, but this is not usually significant due to the overhead of many headers.

*Configuration Areas Affected:*

- Physical and RF Design
- Global CLI Config

## Airtime Optimization

This section focuses on design changes that greatly improve device performance through proper load balancing across APs and channels, leveraging the airtime fairness feature of Aruba's ARM technology, limiting unnecessary traffic, employing multicast rate optimization, and other techniques.

### General Airtime Optimization for All HMD Classes

This section considers wireless load balancing, restricting broadcast and multicast traffic, and disabling "chatty" protocols. These potential configuration changes apply to every type of roaming device.

## Wireless Load Balancing

Highly mobile environments must evenly distribute client load across APs and minimize the duration of client-AP transmissions to get the best efficiency from each radio channel. However, the current 802.11 standard leaves most of this decision making to the client, which lacks both the network-wide perspective and processing intelligence to do the job well. In addition, the myriad of available HMD client radios, drivers, supplicants, and operating systems makes it difficult to provide fair access to the RF spectrum according to their bandwidth capability.

Some of the key decisions that are left to the client include:

- Choosing the most optimal RF band
- Choosing the most optimal 802.11 channel
- Choosing the best access point
- Staying at the highest available 802.11 data rate
- Saving battery life


Aruba offers new infrastructure-based controls that put the WLAN back in the driver's seat to deliver a consistent experience to HMDs. ARM is a collection of features that automatically steer HMD clients to the most efficient channels, prevent slower clients from compromising network performance of faster clients, mitigate co-channel interference to provide the best WLAN performance, and optimize the WLAN by load balancing clients across the spectrum of all APs in the RF neighborhood. Because ARM is high throughput (HT) aware in all 802.11n APs, it will choose the best 40 MHz channel pair (primary and secondary channel) to simplify the integration into a greenfield or legacy installed (20 MHz channel wide) deployment. Best of all, ARM does not require any modification of the roaming device.

In a factory default configuration, the controller has most ARM parameters disabled except for Coordinated Access. The following material describes the features that should be enabled in most customer environments with roaming devices.

### Band Steering: Load Balancing Between Bands

Most dual-band clients will attempt to connect to the first BSSID that responds to its 802.11 Probe Request, which may be on a lower performing PHY type (for example, 802.11b/g). The newer dual-band clients scan (send 802.11 Probe Requests for its configured SSID) on all channels on each PHY type and attempt to connect to the BSSID with the strongest signal. Band steering actively guides 5 GHz-capable HMD clients such as laptops, hand-held scanners, WOWs, patient monitoring devices, and voice handsets to the best available wireless channel by automatically directing it to the 5 GHz band. The result is better noise immunity, fewer sources of interference, and more available channels. See [Figure 3](#).

**Figure 3** *Band Steering*



Here's how it works. Band steering exploits the local probe-response feature of the 802.11 standard. As soon as any Aruba radio sees a probe request in 5 GHz, the device MAC address is placed in a table of 5 GHz-capable clients. The table is automatically propagated to all APs associated to that controller.



Once detected, 5 GHz clients will not be allowed to associate in the 2.4 GHz band except under certain conditions. Band steering also works on Aruba Remote Access Points (APs).

Band steering has the following design considerations:

- Band steering is disabled by default
- You must enable local probe-response



---

Enable probe-response during a maintenance window, as this function causes the VAPs to be re-bootstrapped, which in turn causes a temporary disconnect to devices associated to the SSID.


---

- The band steering table is kept unique on each local controller:
  - There is a table limit of 4,000 entries
  - Entries do not expire
  - Tables are not synced between controllers
  - If an AP reboots, all table entries (up to 4,000) are pushed to the AP

### Spectrum Load Balancing: Load Balancing Across Channels

This feature, depicted in [Figure 4](#), enables APs and controllers to dynamically shift Wi-Fi clients to APs within the same RF neighborhood on underutilized channels. This technique is intended to prevent degraded network performance due to over-subscription. This technique also highly benefits HMD clients in dense office spaces, conference rooms, lecture halls and classrooms, and environments that have high bandwidth applications, because client density is then dynamically balanced among APs in the same vicinity.

**Figure 4** *Spectrum Load Balancing*



The Spectrum Load Balancing feature works in the following manner:

- Distribution of client density across available channels is done by computing the client density metric of each AP within the RF neighborhood.
- When load balancing is active, new Wi-Fi clients that attempt to associate on a fully-subscribed AP will have their associations rejected with a “resource constrained” reason code.

- If clients are sticky and they immediately try to re-connect after their association is rejected, the AP accepts the association.

Spectrum Load Balancing has the following design considerations:

- Spectrum load balancing is disabled by default.
- This feature works best when APs have at least 50% overlap of hearing similar clients.
- The administrator should not expect a constant, even spread of clients across each AP as there are variables that could make the client count differ somewhat.
- The load balancing algorithm does not consider client throughput utilization. However, Aruba legacy load balancing does.
- Local probe-response must be enabled.
- Fast Roaming should be disabled.
- Spectrum load balancing works best on a deployment with the same type of APs.
  - Load balancing does not differentiate 40 MHz HT clients or 20 MHz legacy clients, so a 40 MHz HT client might get load balanced to a legacy AP (802.11abg).
- Only new clients attempting to associate are load balanced.

### **Mode-Aware ARM**

Many HMDs manufactured prior to 2006 have difficulty making good roaming decisions in a dense AP deployment. One way to help the clients is to dynamically reduce the number of APs to match the current load in the environment.

The Aruba Mode-Aware feature dynamically shifts APs with excess capacity in the same RF neighborhood to become temporary Air Monitors. This feature helps all HMD clients quickly find the best available AP during roams and helps clients maintain the best performance possible.

The Mode-Aware algorithm is aware of the physical geography of the network, so it will only disable non-edge APs into temporary Air Monitors (APM) when there is excessive RF coverage.

Co-channel interference mitigation considerations:

- Mode-Aware ARM is disabled by default.
- This feature works in conjunction with coordinated access to a single channel that has no configurable parameters.
- APs cannot be individually configured for Mode-Aware; the feature works across the entire physical AP pool in each AP Group.

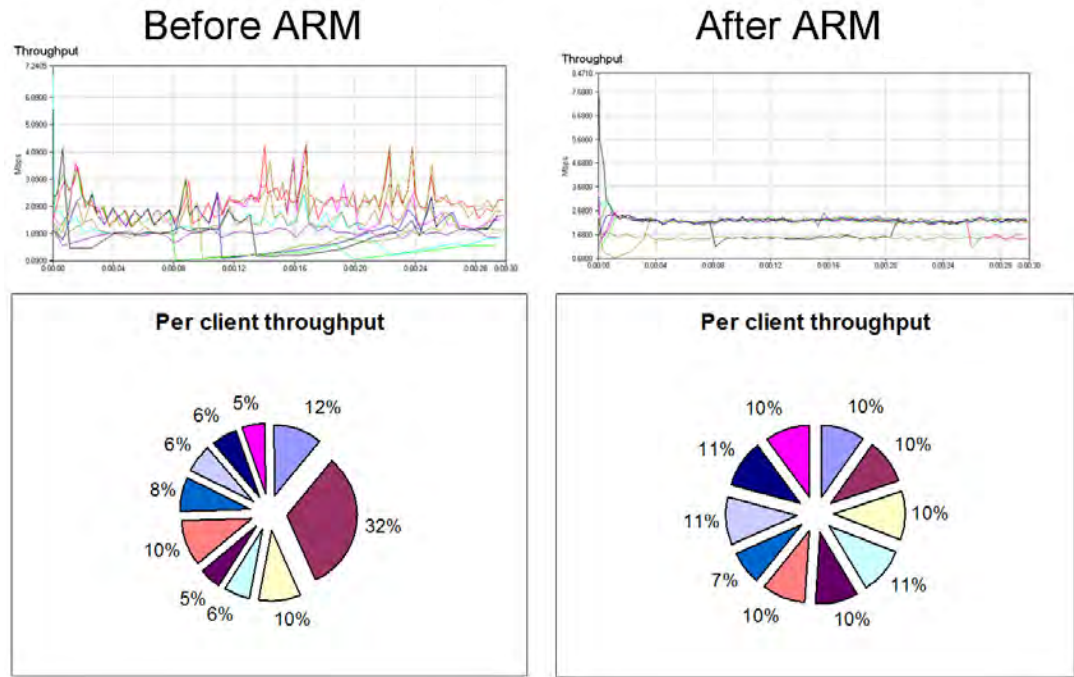
### **Airtime Fairness and Performance Protection**

This ARM feature provides traffic scheduling for dense deployments to deliver equal access to all 2.4 GHz and 5 GHz clients, regardless of wireless chipset manufacturer or device operating system. This feature can also be used to prevent higher speed clients like 802.11n from being compromised by slower speed 802.11b/g clients. This functionality benefits environments where most HMD clients have 802.11a/g/n capabilities and tend to use higher-bandwidth applications.



Airtime Fairness is selectable by radio, and does not require significant resources in the mobility controller CPU. [Figure 5](#) shows the results of a test network with ten different clients before and after enabling Airtime Fairness.

**Figure 5** Airtime Fairness Feature of Aruba's ARM



Airtime Fairness has several design considerations:

- Best to use “Preferred Access” in the Traffic Management profile to allow clients to get their expected throughput based on their supported PHY type
- Works best with client TCP traffic
- Frames with QoS take higher priority
- 802.1X frames will not be shaped

### Wireless Load Balancing Summary

As can be seen from the preceding discussion, comprehensive wireless load balancing requires a number of related features working together. The Aruba implementation of ARM essentially forces the client to make decisions that the infrastructure wants it to. No software or configuration on the client is required to use any of these features.

*Configuration Areas Affected:*

- Global CLI Config
- Virtual AP Profile
- RF ARM Profile

### Restricting Broadcasts and Multicasts

Unnecessary flooding of broadcast and multicast traffic consumes spectrum and prematurely reduces the battery life of roaming devices by requiring them to wake up in order to receive traffic. For example, waking up every device on a channel for an ARP request is expensive. Aruba offers a feature called “broadcast-filter-arp” to address this problem. This feature converts broadcast ARP requests to unicast requests sent directly to the intended client.

For more comprehensive broadcast/multicast filtering, the feature “broadcast-filter-all” is available. This setting blocks all broadcasts except for DHCP. This feature must be used with the “broadcast-filter-arp” feature in order to avoid blocking all ARPs.

*Configuration Areas Affected:*

- Virtual AP Profile

## Limiting “Chatty” Protocols

As of this writing, most organizations do not have a requirement for IPv6 or other “chatty” protocol connectivity. However, Windows Vista comes with IPv6 enabled by default while the Apple iPhone generates large amounts of multicast traffic. Restricting or eliminating this type of connectivity for their wireless users cuts down unnecessary traffic that is not needed for most day-to-day business applications.

Specifically, Aruba recommends using ACLs and settings on the controller to restrict this traffic as follows:

- Limit what devices can appear in the controller’s user table by specifying exactly what subnets and protocols are allowed for HMD clients through the “validuser” IP access list. The following CLI command can be used: `firewall local-valid users`.
- If IPv6 is not required for HMD client connectivity, it is suggested to block it via Ethernet ACL on each mobility controller interface and user-role, as IPv6 quickly consumes user entries on the controller, it is chatty with multicast by default with some devices, and because it is a good general security best practice to disable any unused network protocols to minimize potential risks.
- If netbios-ns, netbios-dgm, mDNS, UPnP, and SSDP protocols are not required for HMD client application connectivity or services, it is strongly suggested to block them in the HMD user-role. These protocols are quite chatty through device queries or announcements and are mainly used for discovering devices in small networks such as in-home networks. Most devices that support these protocols can easily use DNS instead, which is a more optimal protocol for large, highly mobile networks.
- Prevent HMD clients from accidentally being configured as DHCP Servers by blocking the protocol port “udp 68,” which is used for DHCP server replies. This setting should be applied to every HMD user-role.
- When creating ACLs, use netdestination aliases when several rules have protocols and actions in common with multiple hosts or networks, to simplify firewall policy configuration. The netdestination alias allows adding IP addresses by host, network, range, or by using the invert feature. It is best to use “network” to specify a range of hosts when creating a netdestination alias to minimize the number of ACL entries created on the controllers, which have a maximum limit of 4,000 entries.
- Disable spanning-tree in order to not conflict with the uplink distribution switches that usually have the standard spanning-tree or other versions of this protocol enabled in order to provide sufficient redundancy.

*Configuration Areas Affected:*

- Global CLI Config

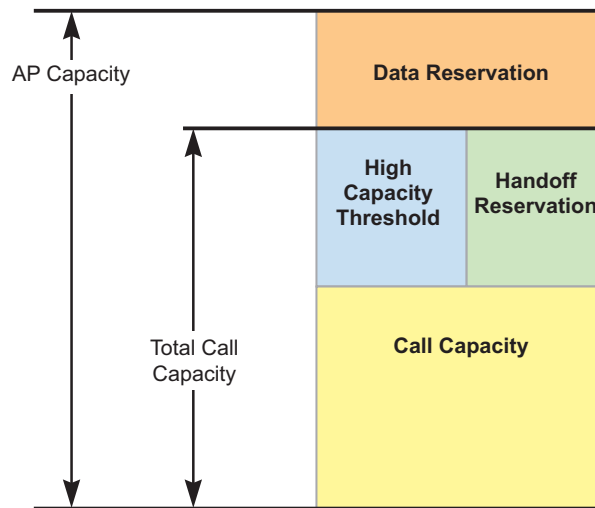
## Voice-Specific Device Configuration

This section considers the planning requirement to ensure sufficient AP capacity for the intended device population prior to deploying voice.

## Capacity Planning

The 802.11 wireless networking protocols are half-duplex and use a contention-based algorithm. As a result, there is a limit to the optimal number of voice clients per AP, depending on the overhead of the VoIP protocol headers, packet sizes, and the encryption used.

**Figure 6** AP Capacity Planning for Call Admission Control (CAC)



Retail\_157

Figure 6 illustrates the potential call capacity of an AP. The top horizontal line in the figure represents the total gross call capacity of the AP and the bottom horizontal line represents a call capacity of zero. The gross call capacity of the AP is diminished by the following areas shown in the figure:

- Data Reservation (top of figure). This amount of the gross call capacity is reserved for data applications. Subtracting out the data reservation leaves the total voice call capacity (labeled Total Call Capacity in the figure).
- High Capacity Threshold and Handoff Reservation. These two shared areas on the diagram further diminish the call capacity. The High Capacity Threshold area is the amount of capacity reserved for peak calling activity so that calls are not dropped during high call demand periods. The Handoff Reservation area is the amount of capacity kept on standby for roaming users who are coming from one AP to another AP.

The net resulting average call capacity of the AP is labeled Call Capacity in the figure.

The recommended maximum per-AP call capacities for clients using G.711 and SIP are listed in Table 5. These values do not indicate the maximum capacity numbers. Rather, these values are the number of calls that can be sustained by the AP with an acceptable background traffic bandwidth.

**Table 5** Call Capacities for Clients Using G.711 and SIP

802.11 Variant	Simultaneous Call Range
802.11b	12 calls
802.11g	25-30 calls if there are no 802.11b clients
802.11b/g	18-20 calls in a mixed 802.11b/g environment
802.11a	20-25 calls

Call capacity with codecs such as G.729 yield up to a 20% improvement over the call capacities just listed. G.711 is widely supported. Because it marginally improves voice quality, G.711 is the recommended choice for VoFi.

Enabling Aruba Call Admission Control (CAC) on an AP helps ensure that the AP is not overwhelmed by simultaneous calls beyond a specified capacity. CAC is aware of the call status of the client (the on-hook/on-call status), which allows the algorithm to make intelligent call balancing and capacity control decisions gracefully with minimal impact to the call quality.

Aruba strongly recommends enabling CAC for production voice deployments. The maximum number of calls supported per AP is a configurable parameter and should be set depending on the background traffic bandwidth required on the AP on the same band as the voice clients. CAC is implemented on a per AP, per radio basis. Set the handoff reservations and the high capacity threshold value to 20%.

These call-based CAC settings are recommended for a single controller environment only. CAC also supports TSpec based bandwidth reservation for voice clients allowing voice clients that don't support TSpec to coexist with the clients that do. TSpec-based CAC can be enabled in a single controller environment in addition to the call based CAC for handsets that support TSpec.

In a multi-controller environment, CAC can be enforced on clients that roam from one controller to another if and only if the clients support TSpec signaling and TSpec signaling is enabled. The recommended setting in a multi-controller environment is to enable both call based CAC for intra-controller CAC enforcement, and TSpec based CAC for both intra- and inter-controller CAC enforcement. The TSpec based CAC enforcement for an inter-controller environment is available as of Aruba OS versions 3.2.

Lab tests show limits on the order of 12-15 calls for 802.11b. However, in practical networks a figure of 10 calls per AP is a reasonable maximum for 802.11b. This figure leaves some bandwidth "reserved" for data traffic and allows for handsets connecting at lower data rates than 11 Mbps.

If 802.11g or 802.11a handsets are used, the situation is much easier: more than 30 simultaneous calls on an AP are possible with these handsets. Lab tests show that up to 76 calls can be handled when all clients connect at 54 Mbps.

A quick capacity check shows whether these figures are reasonable for the expected client density. Take a Wi-Fi cell on the floor plan, estimate the peak number of voice clients (for example, for a cell covering 12 offices/cubicles, perhaps 24 voice clients peak), and then take an estimate of the peak active calls (for example, if 33% of 24 voice clients are on-call = 8 calls peak). Be sure this number does not greatly exceed the capacity of the cell (10 calls for 802.11b). Remember that the CAC algorithm limits calls and maintains quality if thresholds are exceeded, and that load-balancing to adjacent cells usually accommodates excess traffic.

*Configuration Areas Affected:*

- Global CLI Config

## Roaming Optimization

Of the many functions required to achieve good HMD performance, inter-access point handover presents the most challenges. In the space of less than 100 milliseconds, a client should decide it is time to hand over, choose, and associate with the target AP while rigorously maintaining security, and then re-establish any open application sessions. In this section, we consider optimizations to a base design that will greatly enhance the speed and reliability of roaming events.

### General Roaming Optimization for All HMD Classes

This section considers optimal RF signal coverage, VLAN pooling, and fast roaming technologies. These potential configuration changes apply to every type of roaming device.

## Wi-Fi Coverage

Although nearly all currently-available Wi-Fi phones support only 2.4 GHz radios (802.11b and 802.11g), the infrastructure should support 802.11a/b/g everywhere, as the 5 GHz band (802.11a) has many more available channels and is less prone to interference.

Wi-Fi coverage for voice requires more planning than for data users, due to the importance of avoiding dead spots. This is a significant concern because people are inclined to attempt voice calls from places where they would never think of using their PC. The best practice today is to upload a floor plan to Aruba's RF Plan or VisualRF planning tool to identify AP locations, and then perform a "walk-around" site survey to identify any special situations such as large metal objects, thick reinforced concrete walls, and so on.

For voice, it is important that there be continuous coverage, but the APs should not be too close together. Very closely-spaced APs result in extra handover events and can make it more difficult for the client to make a good handover decision. The usual parameters for a planning tool would be for a minimum data rate of 6 Mbps (802.11g) with a 50% overlap between cells (in the Aruba planning tool, set overlap to 150%). As a general rule of thumb, it is best to have a minimum RF coverage of -67 dBm and not higher.

AP spacing will generally be about 20-25 meters (60-75 feet) for data-only networks and 15-20 meters (45-60 feet) where voice is used. Modern WLANs automatically adjust AP transmit power levels and channels for optimum coverage once the APs are installed, so the objective should be to install APs more densely than would be necessary if all APs were running at maximum power (consistent with the 6 Mbps minimum data rates suggested above).

*Configuration Areas Affected:*

- Physical and RF Design

## VLAN Pooling

Use VLAN pools in the virtual AP profile for large networks that require more than one subnet for HMD clients within a specific floor or building. Doing so restricts the size of the broadcast domain, thereby limiting unnecessary traffic.

- Keep each VLAN subnet within a VLAN pool to a 24-bit subnet mask.
- Do not have more than 10 VLANs within a pool so that broadcast or multicast traffic does not consume too much air time access.
- If HMD clients will roam across APs that are bound to different VLANs for the same SSID, then enable VLAN mobility unless the caveats mentioned in the IP Mobility Configuration section require you to enable L3 mobility instead.

### Advantages of VLAN Pooling for Voice Devices

If the voice client supports layer 3 communications between the server and the handset and the number of devices exceeds 200, it is recommended to use VLAN pooling to load balance all the devices associated with the SSID across a number of dedicated voice VLANs. Alternatively, the voice devices can also co-exist with the data devices in the data VLAN provided that the devices and the VLANs are secured.

*Configuration Areas Affected:*

- Virtual AP Profile

## Fast-Roaming Technologies (802.11r and OKC)

Roaming from AP to AP is a function of both the infrastructure and the client. The target figure for interruptions in sensitive applications like a voice call due to handover is around 50 ms. At these levels, the voice HMD will not notice the interruption.



---

Inter-controller roaming is not supported with OKC.

---

Security determines handover performance. Pre-shared keys allow much faster handover than 802.1X environments, although of course they are not as secure. What most organizations want is the security of 802.1X with the speed of a pre-shared key. Two new technologies seek to provide this:

- The ideal balance of security and performance today is when WPA2 is used with 802.1X and opportunistic key caching (OKC) in an Aruba centralized-controller WLAN. Aruba strongly recommends that customers purchase HMDs that support OKC if they intend to use 802.1X authentication for those devices.
- 802.11r is a new standard, intended to improve handover performance. In centralized WLAN architectures such as that provided by Aruba, it offers only limited improvements, particularly if OKC is used (see above). But 802.11r support should be on the checklist for the WLAN vendor.

### Advantages of OKC for All HMDs

Implementing OKC to reduce the time taken to execute a handover eases a number of difficulties:

- An attempted handover is more likely to complete successfully. With full 802.1X re-authentication, several handovers may fail to complete because a poor choice of target AP was made. After several hundred milliseconds elapsed, its signal strength could have dropped significantly, forcing a new handover attempt.
- Because OKC involves only one-fourth the number of frames, it is much less susceptible to frame errors and retries, or to interference from other Wi-Fi and non-Wi-Fi devices. Whereas these effects stretched some full re-authentication events over several seconds, OKC handovers invariably completed promptly.
- There is no penalty in OKC for a client proffering an incorrect Pairwise Master Key (PMK), either because it mistakes the network or AP, or its PMK cache is out of sync with the WLAN. If the authenticator does not recognize the PMK, it ignores it and proceeds with a full 802.1X re-authentication in the usual way.

### Advantages of OKC for Voice Devices

There are also implications for the VoIP layer. Most VoIP clients transmit and receive one frame every 20 ms, and although codecs can now mask single-frame or dual-frame errors, they work best when fed with an uninterrupted stream of frames. A handover interruption on the order of 40 ms normally results in the loss of one or two frames—imperceptible to the listener—rather than the 30 to 50 frames that would be lost during full re-authentication.


Further, VoIP devices maintain jitter buffers to smooth buffering fluctuations and jitter in the network, effectively turning jitter into fixed delay. Modern VoIP clients use adaptive jitter buffers on the order of 80-120 ms which, under normal conditions, maintains a fill level of 50-60 ms. If delay time suddenly increases, as in a handover event, the jitter buffer continues to play out buffered traffic until it expires. At this point, it will take 20-40 ms to recharge when the media stream is restarted. If handover interruptions can be kept infrequent enough that the jitter buffer does not expire, the codec will find it much easier to maintain a constant media stream to the listener, and this is indeed usual with the sub-50 ms interruptions observed during OKC handovers.

Table 6 and the graph in Figure 7 demonstrate some key points regarding handover performance. PMK caching significantly reduces the time taken for handovers, from the order of several hundred milliseconds to less than 75 milliseconds. OKC allows this advantage to be enjoyed for all handover events, not just the special conditions required for PMK caching.

**Table 6** Cisco 7921g Roaming Performance (WPA2-AES Enterprise, EAP-PEAP)

	PMK Cached			
	Intra-Controller Intra-VLAN	Intra-Controller Inter-VLAN	Inter-Controller Intra-VLAN	Inter-Controller Inter-VLAN
<b>Min.</b>	56 ms	62 ms	63 ms	56 ms
<b>Max.</b>	86 ms	67 ms	70 ms	62 ms
<b>Avg.</b>	68 ms	64 ms	66 ms	59 ms

**Figure 7** Nokia Dual Mode Handset Handover Performance



The new Nokia E-series handsets (E65, E51, E66, and E71 in the graph) implement OKC from the IEEE 802.11 standard, and Aruba WLANs have offered OKC for some time. Without any special tuning, this combination of infrastructure and clients demonstrated handover performance previously unseen even with proprietary handover acceleration protocols.

### Design Considerations for OKC and PMK Caching

- OKC is enabled by default in the controller. It should be disabled if no devices in the environment support OKC.
- Aruba recommends enabling “validate pmkid”.

*Configuration Areas Affected:*

- AAA Profile



## Single-Purpose Device Specific Configuration

This section considers special optimizations for single-purpose HMDs in the area of roaming. Aruba recommends two changes over and above the other techniques described earlier:

- Increase the ARM Minimum Transmit Power to 18. This setting helps HMD clients with very sensitive roaming algorithms that require very strong AP signal strength. Apply this setting to the RF ARM profile in both 802.11a and 802.11g radio profiles. By default, the same ARM profile is tied to both radio profiles.
- To provide better interoperability with older HMD clients with controllers that have not been upgraded to ArubaOS 3.3.2.12, it is suggested to remove the highest OFDM 802.11 data rates. Typically, this means 36Mbps, 48Mbps, and 54Mbps. This change is made in the SSID Profile.



---

Be sure to apply this setting to all single-purpose HMD clients as well as the controller.

---

### *Configuration Areas Affected:*

- SSID Profile
- RF ARM Profile

## Voice-Specific Device Configuration

This section considers additional optimizations for Voice HMDs. These changes are in addition to the changes for all HMDs, as well as the single-purpose HMD improvements just mentioned.

### Enable ARM with Voice-Aware and Min/Max Output Power

ARM should be enabled for channel and power management. If this feature is not already configured, Aruba strongly recommends this setting for voice deployments.

Once ARM is running, enable voice-aware scanning. Voice-aware scanning allows the Aruba system to postpone scanning functions on a per-AP basis when it detects an active call on the AP.

In addition to enabling these features, Aruba recommends limiting the minimum and maximum transmit power settings that ARM can use. It is important to match client and AP power. This is especially important for voice devices that typically have limited batteries and radio output.

**Table 7** ARM Transmit Power Settings

Minimum TX Power	Maximum TX Power	Recommendations
0	12	Make sure that the difference between the max and min TX power is no larger than two levels.
12	18	
15	20	
18	30	



## Configure Max Retries, Max Transmit Failures, and Disable Probe Retries

Aruba recommends three additional changes for voice:

- A general best practice for voice deployments is to set the retries on the controller and handset to 2. Because VoIP is delay sensitive, after the packet is delayed, retrying in order to successfully transmit a packet may just add to the latency in the network.



In noisy environments (from the radio point of view), it is recommended to increase the retries value, especially for signalization safety.

- Set the max-tx-fail retries value to 25. This value is the number of consecutive transmitted frames from the AP that are not acknowledged by the HMD client that the frames are destined to. If the station ACKs any frame, then the counter is reset. If max-tx-fail retries value is reached, then the client will be *deauthed* from the AP.
- Probe retries should be disabled

*Configuration Areas Affected:*

- RF ARM Profile
- SSID Profile

## IP Mobility Configuration


The following section describes two common options that customers may use to provide seamless roaming for HMD devices. One option is L2 mobility (see [Figure 8](#)) and the other option is L3 mobility (see [Figure 9](#)).

### General IP Mobility Configuration for All HMD Classes

This section explains the difference between L2 and L3 mobility, and explains how to choose between the two for a given set of client or roaming times requirements. These potential configuration changes apply to every type of roaming device.

With L2 roaming, the roaming user can maintain application connectivity within the roaming domain as long as its layer 3 network address is maintained (does not change). In an L2 Mobility design, the network is designed such that the client maintains its IP address as it roams across controllers and is always assigned an address from the same IP subnet irrespective of the controller or AP it associates to. A general rule of thumb is to limit devices per subnet to 253. However, this number can vary depending on the protocol used and the amount of broadcast or multicast traffic the protocol generates.

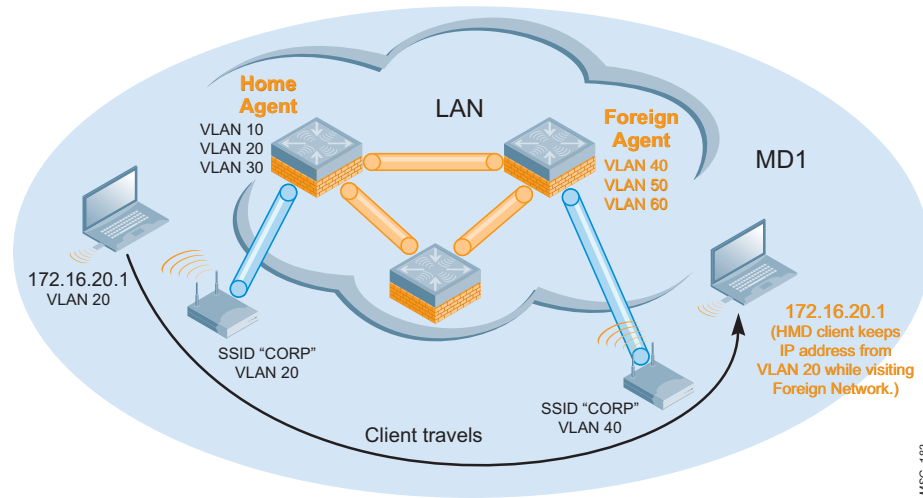
**Figure 8** L2 Mobility



MSG\_187

With L3 roaming, the user is roaming from an AP on Subnet A to an AP on Subnet B. As a result, the layer 3 network address must change in order to maintain layer 3 connectivity on Subnet B. Aruba L3 Mobility allows the HMD client to maintain the same IP address even though it is roaming to a different subnet.

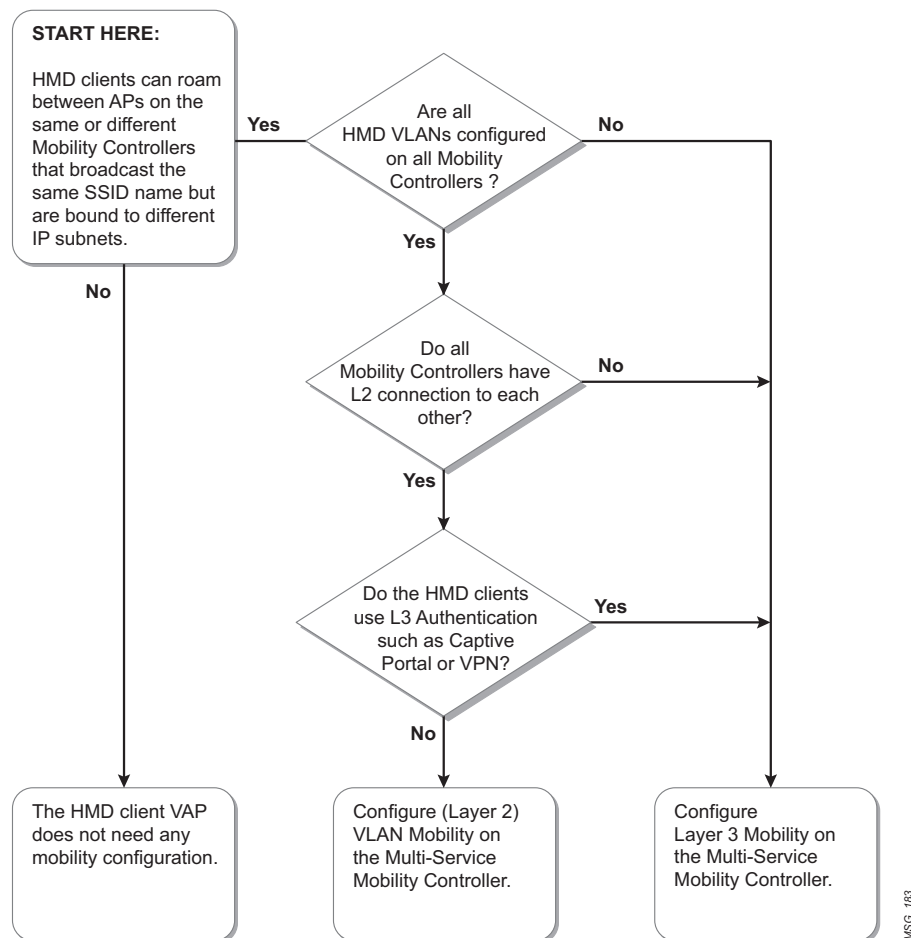
**Figure 9 L3 Mobility**



## Choosing Between Layer 2 and Layer 3 Mobility

Use the chart in [Figure 10](#) to determine if you should be using L2 or L3 mobility.

**Figure 10 IP Mobility Decision Tree**



## Layer 2 (VLAN) Mobility

After VLAN mobility is enabled on a virtual AP, any future association to the APs sharing that profile causes the mobility controller to look in its bridge table to determine if the station is already assigned a VLAN. If an entry is found, the station's VLAN membership is preserved and the BSSID is added to that VLAN. If no bridge entry is found, the station is assigned to the default VLAN of the BSSID or VLAN pool, if any.

The bridge table is aged slowly in order to preserve the VLAN membership. All controllers that participate in VLAN mobility must be members of the same VLANs (at layer 2). Use static generic routing encapsulation (GRE) tunnels if there are topological restrictions.

In an example network offering L2 mobility, you could have two controllers managing different sets of APs. Access point AP1 is connected to controller C1 while access point AP2 is connected to controller C2. The two APs might reside in different VLANs (note that these VLANs are completely independent of the wireless user VLANs). When a wireless client roams from one AP on one controller to a second AP on a second controller, the second controller ensures that user's VLAN assignment is maintained. This ensures that the client retains the same IP address. Session persistence is maintained without any additional control overhead or any inter-controller communication or data redirection.

### Layer 2 (VLAN) Mobility Design Considerations

- All VLANs that require mobility need to be present on all controllers.
- Having more users on different VLANs on an AP increases the overhead for broadcast and multicast traffic on those APs.
- VLAN mobility does not work with layer 3 authentication tools such as Captive Portal or VPN.
- HMD client entries are cached in all controllers that the user visits.
- HMD client L3 session state is not synced between controllers.

An example voice HMD protocol that doesn't require session awareness is Polycom's SVP.

- If machine authentication is enforced with 802.1X profile configuration and it is possible for HMDs to roam between controllers, make sure there is no "aaa internal-userdb use-local-switch" configured on those controllers.
- Even if your primary roaming decision is to use L2 mobility, there may be use cases that require users to be routed from one set of VLANs on Mobility Cluster "X" to an entirely different set of VLANs on Mobility Cluster "Y". In this case, in order to provide seamless roaming, you need to have both L2 Mobility and L3 Mobility configured on the same controller cluster. L2 mobility only allows users to roam between identical sets of VLANs
- Aruba has measured layer-2 roaming times for HMD data clients in the ranges shown below. Exact performance will vary based on device processor speed, radio type, network utilization and other factors.

	With OKC	No OKC	
	Intra-Controller	Intra-Controller	Inter-Controller
<b>Layer 2</b>	19.3 ms	76.5 ms	83.0 ms


## Layer 3 (IP) Mobility

The Aruba Layer 3 Mobility solution is based on the Mobile IP protocol standard, as described in RFC 3344, “IP Mobility Support for IPv4.” This standard addresses HMD clients that need application session connectivity while being mobile within the work environment.

Figure 11 shows the routing of traffic from the Home Agent host to the Foreign agent host when the client is away from its home network. The client’s care-of address is the IP address of the controller in the foreign network. The numbers in the figure correspond to the following descriptions:

1. Traffic to roaming client arrives at the HMD client’s home network via standard IP routing mechanisms.
2. The traffic is intercepted by the home agent in the HMD client’s home network and forwarded to the care-of address in the foreign network via an IP-IP tunnel.
3. The foreign agent delivers traffic to the roaming client.
4. Traffic sent by mobile node is also tunneled back to the home agent.

**Figure 11** Routing of Traffic to Mobile Client within Mobility Domain



Unlike other layer 3 mobility solutions, an Aruba mobility solution does not require the HMD client to install mobility software or perform additional configuration. The controllers perform all functions that enable clients to roam within the mobility domain. For detailed additional information on IP Mobility, please see the *ArubaOS User’s Guide*.

## Layer 3 (IP) Mobility Design Considerations



NOTE

There is no ability to provide L3 mobility to Aruba XSec HMD clients.

### When to Use L3 Mobility

- Use L3 mobility if the HMD client’s IP address is required not to change when the client crosses layer 3 boundaries from its home network.
- Use L3 mobility if the HMD client’s layer 3 state, role, or session, established through Captive Portal or Virtual Private Network, needs to be preserved during inter-mobility controller roaming. An example voice HMD protocols that require session state are SIP, NOE, and SCCP.

## Configuring L3 Mobility

- If your primary roaming decision is to use L3 mobility in order to route users to different sets of VLANs, there is no need at all to also have L2 mobility configured on the same master/local controller cluster. The reason for this is that L3 mobility handles all roaming cases.
- The configuration parameter “router mobile” must be enabled on all controllers that will need L3 mobility.
- Every virtual AP that requires L3 mobility should also have mobile IP enabled.
- Every HMD client VLAN that requires L3 mobility must have an IP interface configured.
- If the same HMD VLAN exists across more than one controller, the VLAN(s) should be added as a potential HA in the home agent table. This helps prevent one controller from getting overloaded with all the other controllers’ clients that share the same VLAN.
- For voice HMDs and single-purpose HMDs such as scanners, printers, and patient monitors that perform silent roaming to maximize their battery life, you should enable the HA Discovery on-association parameter (retail customers may use this for their equipment [scanners, printers, or any wireless equipment with static IP addresses] and have mentioned the best practice to still enable "on association" in mobility for that purpose.) on each virtual AP that needs L3 mobility. This parameter is useful when a controller does not have a state and the HMD client performs an inter-controller move. When the system is notified of an 802.11 association and the on-association parameter is enabled, the controller performs HA discovery based on the HMD client MAC address. Because this HMD client should exist somewhere already, the new controller can become an FA. If no controller claims to have a session for this HMD client, the HA discovery fails and this controller becomes the HA.
- You should also enable the IP mobile proxy on-association parameter. This is useful for an intra-controller move when a controller already has a state for the HMD client and it performs an AP1 to AP2 move. The moment the HMD client moves to AP2 with this feature enabled, the controller detects this move as soon as it receives an 802.11 association message.
- If Virtual Router Redundancy Protocol (VRRP) is being used on the controllers for providing redundancy to APs, each of the mobility domain’s home agent tables should be configured with the LMS VRIP as the Home Agent IP instead of the Controller Switch IP (this feature is only supported in ArubaOS 3.3.2.x and above).
- A mobility domain needs to be made active and have a home agent table. The home agent table must be configured with the list of every subnet, mask, and home agent address that is valid for one mobility domain. There must be an entry in the home agent table for every VLAN on every switch on which stations are expected to appear. If more than one controller is supposed to provide mobile IP service on a given VLAN, one entry for each controller is needed. The specified HA must have IP connectivity to all Controllers within the configured mobility domain or else there will be a delay in station roaming as specified in the HA failure bullet below. Protocol UDP port 434 must be allowed among all controllers configured within a mobility domain.
- Make sure that mobile IP proxy event thresholds is set at 10 (events per second) for best scalability.

## Architecture of L3 Mobility

- Controllers can belong to more than one mobility domain.
- A mobility domain should contain all controllers that a HMD client could potentially roam among without dropping its WLAN link due to lack of RF coverage. For example, if an HMD client can roam to all buildings in a campus without losing any WLAN connectivity, the controllers that service those buildings should belong to the same mobility domain.
- If the controllers have virtual APs configured to use VLAN pooling, it is suggested but not required that the VLAN pools be kept different on each controller that has L3 mobility enabled in that specific VAP.
- The default mobility domain cannot be deleted and is used for any controller that has mobile IP enabled but is not specifically bound to any mobility domain.

### HA Failures

- In the case of an HA failure while the client is visiting an FA and is actively being homed back to its HA, the FA will start owning the client once the mobility process can detect that the client is using a stale IP address. It determines this by discovering there is no available HA for its IP subnet in the home agent table of that mobility domain because the HA is down. The FA recovers this client by sending 802.11 deauth messages (total of 3) at 30-second intervals until it renews its IP address so the FA can become the HA for the client. In this instance, the client will have a new IP address.

### Firewall or Access Policy

- The HMD client's data traffic will be screened at its HA due to a firewall or access policy.

### Multicast Streams

- A wireless client with an active multicast stream will experience an interruption when roaming between HA and FA because the mobility client IP address does not change when roaming to the FA, henceforth the upper layer protocol does not detect any change in point of attachment and no new IGMP joins are generated.

### Roaming Times

- Aruba has measured layer-3 roaming times for HMD data clients in the ranges shown below. Exact performance will vary based on device processor speed, radio type, network utilization and other factors.

	With OKC	No OKC	
	Intra-Controller	Intra-Controller	Inter-Controller
<b>Layer 3</b>	17 ms	185 ms	229 ms

*Configuration Areas Affected:*

- Global CLI Config
- Virtual AP Profile

## IP Multicast Optimization

Up to this point, this Guide has emphasized methods of reducing or eliminating multicast traffic to improve roaming device battery life and network performance. However, some devices and applications depend on multicast transport and services must be provided to them.

Some common applications for multicast at this writing include healthcare patient monitoring systems, IP television streaming, and IP security camera streaming. These applications tend to be mission critical and must work reliably at all times.

### General Multicast Optimization for All HMD Classes

This section considers design issues and configuration notes for ensuring optimal multicast performance for devices that require it, while minimizing bandwidth consumption on the LAN and over the air. These potential configuration changes apply to every type of roaming device.

### Multicast Design Considerations

- It is important that controller and wired network infrastructure capacity planning is properly executed before the deployment of multicast streams.
- The controller should NOT be the default gateway for HMD clients, because the controller is not intended to support multicast routing.

- All HMD VLANs should have L2 connectivity to the uplink router that is enabled for multicast routing for those specific subnets.
- In an effort to reduce the amount of traffic replication and processing on controllers and within the wired infrastructure, it is highly recommended to enable IGMP snooping (v1 and v2 for IPv4 and MLD for IPv6) on the controllers. This feature allows the controller to snoop IGMP messages from HMD clients on a per VLAN basis and use this information to create a multicast forwarding table that defines mapping of a particular multicast group and outgoing port/interface(s). This table takes care of forwarding the multicast stream only to the AP tunnels that serve clients who are subscribed to that multicast stream, and prevents unnecessary use of controller and wired switch/router datapath.




---

IGMP snooping does require the VLAN interface to have an IP address.

---

- By default, multicast and broadcast traffic is transmitted at the “lowest 802.11 transmission rate” configured on the SSID. This may not be an acceptable solution in terms of required wireless throughput, because 5 GHz radio is configured with 6 Mbps, and 2.4 GHz radio with 1 Mbps as the lowest transmission rate, by default. Therefore, it is strongly suggested that the APs provide higher performance to HMD clients by forwarding multicast packets in the air at the lowest 802.11 data rate of an associated client by means of the SSID profile parameter “mcast-rate-opt,” which significantly improves HMD client data throughput overall.
- Controllers can use QoS to prioritize and perform ToS assignments to the ingress/egress multicast traffic on the wire by assigning a “session-acl” to the particular Ethernet port and on the HMD client user-role through Access Control List (ACL) configuration policies. Assigning the session-acl helps prioritize the multicast streams within the controller and wired infrastructure datapath.
- For networks that don’t require Multicast Domain Name Service (mDNS) for HMD clients, it is strongly suggested to block this protocol in HMD user-roles because it is very chatty and can consume a large amount of bandwidth with the multicast announcements and queries that get flooded to all VLANs that have mDNS capable devices. Blocking mDNS should not prevent connectivity because normal Domain Name Service (DNS) can be used to discover other network services and devices.
- The controllers can support up to 1024 multicast groups.
- Bandwidth contracts can be applied to HMD VLANs to help unicast traffic remain unaffected in such conditions.
- If required, the controllers can also be instructed to:
  - Limit the set of “user-roles” that can have access to multicast traffic, through the use of its role-based stateful firewall.
  - Assign “time-of-day” to the firewall rules, allowing view of the multicast traffic only during certain time periods during the day.
  - Allow multicast access only at certain groups of APs terminating on the controller. These configuration options may also help in reducing the amount of undesired multicast load within the wireless and wired network infrastructure.
  - For large networks that don’t require HMD clients to use multicast streams, it is strongly suggested to configure bandwidth contracts that limits multicast and broadcast traffic to 1 Mbps for every HMD client VLAN.

*Configuration Areas Affected:*

- Global CLI Config
- SSID Profile



## Interference Resistance

HMDs are more subject to varying RF conditions (and therefore interference) than stationary wireless devices. For that reason, when you design a network with HMDs, it is important to consider that the HMDs will most likely meet adverse RF environments in the course of their usage.

### General Interference Resistance for All HMD Classes

More and more individuals and businesses are buying wireless products that complement their highly mobile lifestyle. Only a few unlicensed frequency bands are available for this purpose. As a result, many of these wireless products share the same RF spectrum as 802.11 devices, particularly the unlicensed Industrial, Scientific, and Medical (ISM) 2.4 GHz band. These devices do not support the 802.11a/b/g/n protocol, yet operate in the same spectrum and thereby become interferers for WLAN clients.

There are four primary types of interference with which roaming devices must contend. Most of these devices use either Direct Sequence Spread Spectrum (DSSS), Frequency-Hopping Spread Spectrum (FHSS), or Fixed Frequency technology. The most common 2.4 GHz non-802.11 devices are microwaves, Bluetooth devices, cordless phones, wireless USB devices, audio/video transmitters/receivers, and wireless gaming devices. Dense 802.11b/g/n AP deployments without advanced RF management technology can also become interferers with co-channel and adjacent channel overlap.

This section explains all four types of interference, and discusses mitigation strategies for HMDs. These potential remedies include both physical device changes as well as controller configuration changes. These approaches apply to every type of roaming device (multi-purpose, single-purpose and voice).

### FHSS and 802.11b/g Co-Existence Design Considerations

In many companies, there are no restrictions enforced on which devices can be brought to work in order to offer convenience to each individual. For example, devices such as a 2.4GHz (non-802.11 device) phone or Bluetooth (non-802.11) device for wireless connectivity causes interference to 802.11 devices. However, some networks that have legacy non-802.11 a/b/g/n devices will eventually be upgraded to the newer technology to offer higher performance to the HMD clients. Such upgrades are done in a certain transition time that requires both technologies to co-exist while all the necessary equipment is fully upgraded in each location.

The characteristics of 802.11 FHSS are:

- Operates in the 2.4 GHz frequency range
- The transmitter carrier frequency changes in a pseudo-random fashion
- The narrowband signal is transmitted at 1 MHz wide in any one of the (US) available channels
- FHSS devices are limited to a dwell time of 400 ms

FHSS devices can significantly reduce 802.11 performance due to the high volumes of Rx interrupts generated when the radio tries to respond to the PHY events by attempting to decode non 802.11 signals. This interference also causes ping losses and packet drops at the AP. For this reason, FHSS devices should be phased out of areas with business-critical 802.11 wireless networks as quickly as possible.

The following are design considerations for networks where FHSS and 802.11b/g must co-exist:

- Follow the guidelines in the applicable Aruba base design (*Campus Wireless Networks Validated Reference Design* or *Retail Wireless Networks VRD*) for best AP placement
- Make sure the FHSS APs are located at least 20 feet from APs
- Verify that the uniform SNR in the environment exceeds the minimum requirement of the HMD manufacturer
- Make sure the controller is using ArubaOS 3.3.2.12 or higher.



- Configure a 1 Mbps bandwidth contract on all HMD VLANs. The purpose of this is to limit the amount of wireless traffic on the air to reduce the incidence of collisions and associated retransmissions.

There should be a different bandwidth contract name for every "entity" (like VLAN # or user-role) so that the Mobility Controller can create a different bandwidth contract # in the datapath for each interface and user-role instead of sharing one bandwidth contract throughput limit across all "entity" that share that one contract name. The following is an example of creating different bandwidth contract names according to the "entity" it will be bound to:

```
config t
  aaa bandwidth-contract vlan1_1Mbps_bandwidth_contract mbits 1
  interface vlan 1
  bandwidth-contract vlan1_1Mbps_bandwidth_contract

  aaa bandwidth-contract vlan2_1Mbps_bandwidth_contract mbits 1
  interface vlan 2
  bandwidth-contract vlan2_1Mbps_bandwidth_contract

  aaa bandwidth-contract vlan3_1Mbps_bandwidth_contract mbits 1
  interface vlan 3
  bandwidth-contract vlan3_1Mbps_bandwidth_contract
end
```

- Enable ARM with these Profile parameters
  - Enable mode aware to allow APs to dynamically switch themselves to temporary Air Monitors when they determine there is too much co-channel overlap, and to provide dynamic self-healing
  - Increase the minimum transmit power to 15 dBm
- SSID Profile parameters
  - Use only 802.11b rates
    - Basic rates: 1, 2
    - Supported rates: 1, 2, 5, 11
  - DTIM period should match the HMD manufacturer best practice setting
  - Do not hide the SSID
  - Do not deny broadcast probe requests
  - Increase maximum retries to 15
  - Enable multicast rate optimization
- Increase Noise Immunity level to 3, 4 or 5 (the default is 2)

ArubaOS 3.3.2.11 adds a new parameter to the *rf dot11g-radio-profile* command in the CLI. This parameter allows the user to set the interference immunity on the 2.4 GHz band. The default setting for this parameter is level 2, which is the same level used by the access point in older releases. When performance begins dropping due to interference in the band, the level can be increased to 5 for improved performance. However, increasing the level makes the AP slightly “deaf” to its surroundings, causing the AP to lose a small amount of range.

The levels for this parameter are:

- Level-0: No ANI adaptation
- Level-1: Noise immunity only
- Level-2: Noise and spur immunity (default)
- Level-3: Level 2 and weak OFDM immunity

- Level-4: Level 3 and FIR immunity
- Level-5: Disable PHY reporting



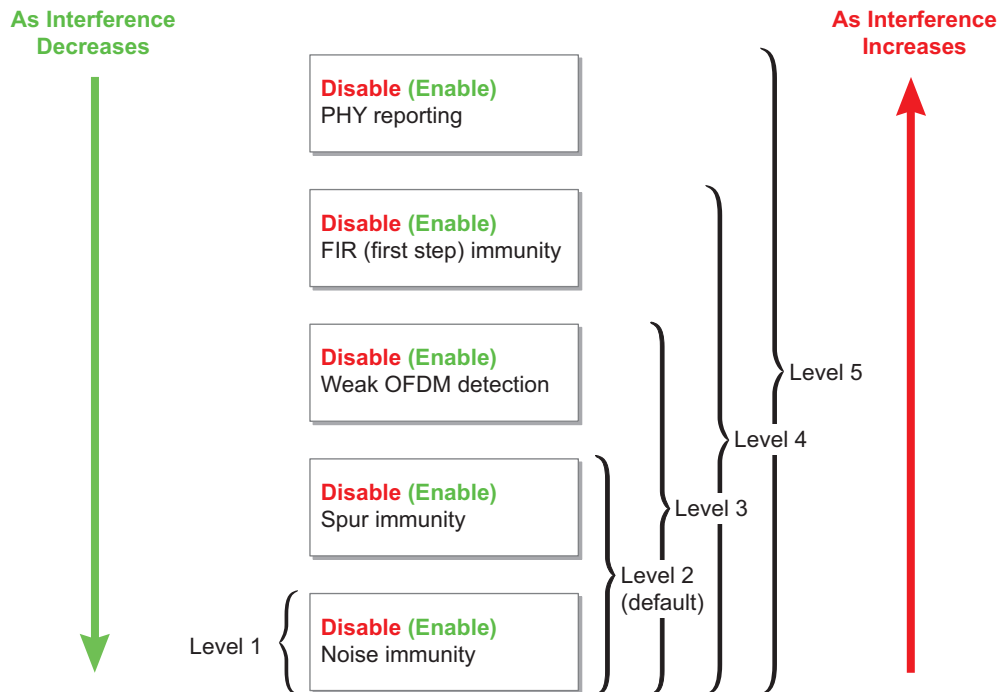

---

If Noise Immunity is enabled, Receive Sensitivity Tuning Based Channel Reuse must be disabled.

---

These settings take advantage of capabilities in the Atheros chipset. In order to mitigate the effects of the non 802.11 interferers, Atheros provides various levels of adaptive noise immunity (ANI). While is used mostly to filter out the effects of board noise, it has been leveraged to adaptively tune the radio sensitivity in the presence of the above interferers.

**Figure 12** *Noise Immunity Levels*



*Configuration Areas Affected:*

- Physical and RF Design Optimizations
- Global CLI Config
- SSID Profile
- RF Radio Profile
- RF ARM Profile

### Fixed Frequency Interference

There are several wireless devices that use fixed 2.4 GHz frequencies, such as security cameras, other analog video transmitters, some cordless phones, and wireless keyboards, all of which interfere with 802.11b/g devices. Follow this series of actions to mitigate fixed-frequency interference sources:

- Locate the device(s) causing interference
- If the device supports manual RF configuration, reduce transmit power or move the device to a non-interfering channel
- Use ARM to dynamically adjust the channel plan around the interference source. ARM contains error and noise floor threshold features to avoid excessive errors or high non-802.11 noise that would greatly disrupt HMD client connectivity and performance.

- Migrate the device to use 802.11. For example, replace 2.4 GHz cordless phones with VoFi handsets, or wireless video cameras with IP cameras.
- Replace the device with one that uses a non-interfering frequency band. This is simple for cordless phones and video cameras, but not possible for microwave ovens.
- Relocate the device
- Purchase 5 GHz-capable WLAN clients and dual-radio APs

By default, the ARM profile has its Error Rate Threshold set at 50%. This means that if the combined percentage from 802.11 PHY and MAC errors on an AP's channel ever reaches 50% (it typically should be less than 20%) for an Error Rate Wait Time of 30 seconds, the AP will move to a different and cleaner channel to provide better connectivity to the HMD clients. The ARM profile also has its Noise Threshold set to -75 dBm by default. This means that if the Noise Floor on an AP's channel ever exceeds this value (it typically should be -90 dBm to -100 dBm) for a Noise Wait Time of 120 seconds, the AP will move to a cleaner channel. Both the Error Rate Threshold and Noise Threshold can be adjusted to be more sensitive to these severe conditions (for example, set Error Rate Threshold to 35% and Noise Threshold to -80 dBm with the default Wait times).

### 802.11 Co-Channel and Adjacent Channel Interference

Shared office buildings with multiple neighboring APs can cause co-channel interference (CCI) and adjacent channel interference (ACI). Co-channel interference is caused when multiple 802.11 APs are physically on the same channel and can hear one another. This can occur on any channel in either the 2.4 GHz or 5 GHz bands.

The network can still function in the presence of CCI, but can suffer reduced performance without any coordinated access support (such as that offered by the ARM feature). The 802.11 PHY layer uses Collision Sense Multiple Access with Collision Avoidance (CSMA/CA). In order for CSMA/CA to transmit, it senses the medium to determine if another station is transmitting. If the medium is busy, the station defers its transmission until the end of the current transmission. After deferral, the station selects a random backoff interval.

ACI is caused by more than one AP being physically on channels that overlap within their channel width (for example, AP1 on channel 1 (which has center frequency of 2412) and AP2 on channel 3 (which has center frequency of 2422)). Make sure that ARM features and their default settings are enabled so that the entire Aruba WLAN can intelligently pick the best performing channels in their current environment.

Follow this series of actions to mitigate CCI and ACI:

- Enable ARM for dynamic channel selection and output power assignment
- Enable Mode-Aware ARM to dynamically convert excess APs into Air Monitors
- Enable RX Sensitivity Tuning Based Channel Reuse

In some dense deployments, it is possible for APs to hear other APs on the same channel. This creates co-channel interference and reduces the overall utilization of the channel in a given area. Channel reuse enables dynamic control over the receive (Rx) sensitivity in order to improve spatial reuse of the channel.

The channel reuse mode is configured through an 802.11a or 802.11g RF management profile. You can configure the channel reuse feature to operate in either of the following three modes; static, dynamic or disable. (This feature is disabled by default.)

- **Static mode:** This mode of operation is a coverage-based adaptation of the Clear Channel Assessment (CCA) thresholds. In the static mode of operation, the CCA is adjusted according to the configured transmission power level on the AP, so as the AP transmit power decreases as the CCA threshold increases, and vice versa.

- **Dynamic mode:** In this mode, the Clear Channel Assessment (CCA) thresholds are based on channel loads, and take into account the location of the associated clients. When you set the Channel Reuse feature to dynamic mode, this feature is automatically enabled when the wireless medium around the AP is busy greater than half the time, and the CCA threshold adjusts to accommodate transmissions between the AP and its most distant associated client.

The following example is a CLI configuration of how to configure Dynamic Receive Sensitivity Tuning:

```
!  
rf dot11g-radio-profile <802.11g rf profile name>  
  channel-reuse <disable, dynamic, or static>  
  channel-reuse-threshold <Rx Sensitivity Threshold value in -dBm>  
!
```

- **Disable mode:** This mode does not support the tuning of the CCA Detect Threshold.



---

If Receive Sensitivity Tuning Based Channel Reuse is enabled, then Noise Immunity must be disabled.

---

The previous chapter was written to help you understand unique mobility-related design principles in order to plan the WLAN so that it readily accommodates highly mobile devices. Guidelines were given on how to set up the WLAN infrastructure and incrementally optimize one of the Aruba published base designs to accommodate HMDs.

The next three chapters now give you specific information on how to implement those design principles on the Aruba Mobility Controllers.

- This chapter applies to all HMDs; the configuration changes listed here should be used for multi-purpose, single-purpose, and voice HMDs.
- [Chapter 5, “Configuring Incremental Settings for Single-Purpose HMDs”](#) provides some enhancements specific to single-purpose HMDs.
- [Chapter 6, “Configuring Incremental Settings for Voice HMDs”](#) provides additional enhancements just for voice devices. These are in addition to the changes in the previous chapters.

The sequence of topics in all of these chapters is the same as presented for [Chapter 3, “Understanding Design Principles for Roaming Devices”](#), with CLI examples provided for each area that touches the controller. You can use the material in this section to extract optimum performance from the network as well as from the HMDs themselves.

This chapter addresses advanced topics in ArubaOS configuration. As such, the reader is expected to have a thorough understanding of profile-based configuration as used in ArubaOS 3.x, and general familiarity with the ArubaOS Profile hierarchy. The reader should also be deeply familiar with the operation and configuration of virtual APs. For further information on these topics, please refer to one of the Aruba published base reference designs, or the *ArubaOS User’s Guide*.



---

We assume that the Aruba Mobility Controller’s factory default settings have already been programmed in accordance with an Aruba base design.

---

## Device Configuration

This section reviews how to implement the design principles related to Device Configuration.

### Configure Optimal HMD Environment Device Settings

Use the five-step process presented in Chapter 4 to obtain the key device-specific default values that should be changed for your particular deployment scenario:

1. Update all devices to the latest available firmware.
2. Make a list of all configurable settings and default values.
3. Contact the device vendor's systems engineer and request a list of the latest known "best practice" settings for your specific deployment scenario.
4. Make the same request of your Aruba systems engineer.
5. Conduct a pilot test to experiment with these values and others that seem relevant but may not be known to the systems engineers.

### Shared or Dedicated SSIDs

Complete the assessment process explained in [Chapter 3, "Understanding Design Principles for Roaming Devices"](#) to determine whether the specific HMDs that will be used in your environment will require one or more dedicated SSIDs. Aruba strongly recommends that the number of SSIDs be kept as low as possible to limit unnecessary LAN and RF overhead. Therefore, the bias should be against a dedicated SSID unless one of the criteria clearly applies.

Dedicated SSIDs are implemented just like any other SSID. Create a Virtual AP with corresponding SSID and AAA Profiles. Also create any supporting profiles required by those profiles or the mobility optimizations required further on in this section.

### Enable 802.11h (Transmit Power Control)

To help provide better connectivity for HMD clients that need to see AP country information, enable 802.11h in both the 802.11a and 802.11g RF profiles. This setting causes the AP to add the 802.11d country and the power constraint information element in the 802.11 header of Beacons and Probe Responses.

1. Enabling 802.11h in the RF profile will turn on the country and power constraint information element in the 802.11 header of all SSIDs in the corresponding AP group.	<pre>! rf dot11a-radio-profile "&lt;802.11a rf profile name&gt;"     dot11h1 ! rf dot11g-radio-profile "&lt;802.11g rf profile name&gt;"     dot11h1 !</pre>
--	--

## Airtime Optimization

The next few pages consider how to implement the design principles related to Airtime Optimization.

### Wireless Load Balancing

- Enable band steering for each virtual AP profile that is broadcasting the SSID on both RF bands (2.4 GHz and 5 GHz) to carefully nudge HMD clients to the highest performing channel.

1. This VAP profile parameter should only be applied to VAPs that are configured for all bands (a/b/g). Local probe response must be enabled as well.	<pre>! wlan ssid-profile &lt;HMD client ssid name&gt;   local-probe-response<sup>1</sup> wlan virtual-ap "&lt;HMD client vap name&gt;"   band-steering !</pre>
---	--

- Enable Spectrum Load Balancing in the RF profile for each PHY type. This helps keep the most optimal client density across all APs by automatically and intelligently assigning HMD clients to the best available channel and AP within the RF neighborhood.

1. Spectrum Load Balancing requires ArubaOS 3.3.2.12. The feature may be enabled on a per-radio basis.	<pre>! rf dot11a-radio-profile "&lt;802.11a rf profile name&gt;"   spectrum-load-balancing<sup>1</sup> ! rf dot11g-radio-profile "&lt;802.11g rf profile name&gt;"   spectrum-load-balancing<sup>1</sup> !</pre>
--	--

- Enable Mode Aware ARM to dynamically convert excess APs into Air Monitors in response to changing load conditions.

1. This ARM profile feature allows the APs to dynamically choose the best coverage pattern of APs and air monitors.	<pre>! rf arm-profile "&lt;arm profile name&gt;"   mode-aware<sup>1</sup> !</pre>
---	---

- Enable Airtime Fairness to allocate access to the channel on a scheduled basis among all clients on an AP.

<p>1. There is no default 'wlan traffic-management-profile'. It must be created.</p> <p>2. Valid policy types are default, fair-access, and preferred-access.</p> <p>3. The 'wlan traffic management profile' is applied separately to each radio at the AP group level.</p>	<pre>! wlan traffic-management-profile "&lt;wtm profile name&gt;"<sup>1</sup>   bw-alloc virtual-ap default share &lt;percentage&gt;   shaping-policy &lt;policy type&gt;<sup>2</sup> ! ap-group "&lt;ap group name&gt;"<sup>3</sup>   dot11a-traffic-mgmt-profile &lt;wtm profile name&gt;   dot11g-traffic-mgmt-profile &lt;wtm profile name&gt; !</pre>
--	--

## Restricting Broadcasts and Multicasts

- Enable the virtual AP profile “broadcast-filter arp” parameter to help limit the flooding of broadcasts sent into the air. This feature converts broadcast ARP requests to unicast requests sent directly to the client. This feature can help improve battery life because HMD clients do not have to wake up from 802.11 power-save sleep state to receive irrelevant ARP packets.
- If no HMD client application requires multicast protocol support then enable the virtual AP profile “broadcast-filter all” feature, which drops all multicast and broadcast traffic in the air. Obviously, this feature should be used with caution. Once enabled, it helps conserve battery life and offers optimization by reducing unnecessary traffic that is usually sent in the air. This feature:
  - a. *Must* be enabled with VAP parameter “broadcast-filter arp” or else it will drop broadcast ARP, as well.
  - b. Does not drop DHCP broadcasts.
  - c. Helps work around some operating systems that send DHCP Discoveries or Requests with the broadcast bit on.

<ol style="list-style-type: none"> <li>1. This VAP profile parameter helps limit the flooding of broadcasts sent into the air by converting broadcast ARP requests to unicast requests which are sent directly to the client.</li> <li>2. This VAP profile parameter should only be enabled on a SSID where HMD clients do not require multicast protocol support because it does drop all broadcast and multicast traffic except for DHCP. If this feature is enabled, then the VAP profile parameter “broadcast-filter arp” <i>must</i> be enabled as well.</li> </ol>	<pre>wlan virtual-ap "&lt;HMD client virtual ap name&gt;"   broadcast-filter arp<sup>1</sup>   broadcast-filter all<sup>2</sup> !</pre>
--	---



## Limiting “Chatty” Protocols

The following configuration suggestions have some parameters meant for customers that do not require IPv6 and other chatty protocol connectivity. Limiting this type of connectivity for their wireless users cuts down unnecessary traffic that is not needed for most day-to-day application and network use. See [Chapter 3, “Understanding Design Principles for Roaming Devices”](#) for a detailed explanation of each setting.

1. If Remote APs are deployed in the network, then create a netdestination name “<Remote_AP_Pool_Name>” that should contain all Remote AP inner IP addresses across all mobility controllers that are terminating Remote APs.	<pre>! netdestination &lt;Remote_AP_Pool_Name&gt;<sup>1</sup> network &lt;network IP address&gt; &lt;subnet mask&gt;</pre>
2. The netdestination name “<HMD_Client_IPs>” should contain all HMD client IP addresses across all mobility controllers.	<pre>! netdestination &lt;HMD_Client_IPs&gt;<sup>2</sup> network &lt;network IP address&gt; &lt;subnet mask&gt;</pre>
3. This valid user acl limits what can be put in the mobility controller’s user-table. By default, everything is allowed. This configuration suggests allowing dhcp, nat-t, l2tp, all Remote AP inner IP addresses, and all HMD client IPs.	<pre>! ip access-list session validuser<sup>3</sup>   any any svc-dhcp permit   any any svc-natt permit   any any svc-l2tp permit   alias &lt;Remote_AP_Pool_Name&gt; any any permit   alias &lt;HMD_Client_IPs&gt; any any permit   any any any deny log</pre>
4. IPv6 is disabled by default, but take this configuration step if it shows enabled with the CLI command “show ipv6 firewall”.	<pre>! no ipv6 firewall enable<sup>4</sup></pre>
5. Prevent any IPv6 traffic passing through the mobility controller by applying this ACL on all controller interfaces and user-roles. Create this Eth ACL on all mobility controllers before applying it to any user-role on the master controller.	<pre>! ip access-list eth no-ipv6-acl<sup>5</sup>   deny 0x86dd   permit any</pre>
6. This ACL prevents any HMD client discovering services and devices through Multicast Domain Name Service. Apply this ACL to all user roles.	<pre>! ip access-list session deny_mDNS_acl<sup>6</sup>   any any udp 5353 deny</pre>

7. This ACL prevents any HMD client discovering services and devices through Universal Plug and Play and Simple Service Discovery Protocol. Apply this ACL to all user-roles.	<pre> ! ip access-list session deny_SSDP_and_UPnP_acl<sup>7</sup>   any host 239.255.255.250 any deny   any host 239.255.255.253 any deny </pre>
8. This ACL prevents any HMD client discovering services and devices through NetBios protocol. This should be applied to all user-roles.	<pre> ! ip access-list session-acl deny_netbios_acl<sup>8</sup>   any any udp 137 deny   any any udp 138 deny </pre>
9. This ACL is in the logon system user-role by default but it should be applied to all user-roles to block any wireless device from acting as a DHCP server.	<pre> ! ip access-list session deny_client_acting_as_server_acl<sup>9</sup>   user any udp 68 deny </pre>
10. This is an example user-role with the suggested protocol deny statements. Notice that the eth-acl should be at the top of the list. Some administrators may not want an “allowall” session acl in the HMD user role, so make sure to create ACLs that specify which protocols are allowed and apply it to the user-role. Note there is an implicit deny at the end.	<pre> ! user-role &lt;wireless user role name&gt;<sup>10</sup>   access-list eth no-ipv6-acl   session-acl deny_mDNS_acl   session-acl deny_SSDP_and_UPnP_acl   session deny_netbios_acl   session-acl deny_client_acting_as_server_acl   session-acl allowall </pre>
11. Disable spanning-tree because there are usually other L2 switches that are configured for it, and might even have a different feature support that would conflict with the mobility controller implementation.	<pre> ! no spanning-tree<sup>11</sup> </pre>
12. The IPv6 eth acl should be applied to all interfaces on all mobility controllers.	<pre> ! interface [all active Fastethernet/ gigabitethernet/port-channel] &lt;slot/port value&gt;<sup>12</sup>   ip access-group no-ipv6-acl in ! </pre>

## Roaming Optimization

In this section, we cover how to implement the design principles related to Roaming Optimization.

### Wi-Fi Coverage

It is essential that the coverage area where HMDs will be operating has been properly RF designed. It is also vital that the installed coverage be measured and found to meet or exceed the expected values for Received Signal Strength Indication (RSSI) and signal-to-noise ratio (SNR). Good coverage is the first step to good roaming performance; conversely, poor coverage is certain to deliver an unsatisfactory user experience.

Aruba base designs contain extensive recommendations for determining the number and placement of APs and optional external antennas. Aruba also offers a variety of RF planning tools such as Aruba RF Plan, AirWave VisualRF, and the Aruba Outdoor 3D Planner. These tools can be used to design coverage that meets the requirements for the specific HMD classes that will be used in your facilities.

### VLAN Pooling

- Use VLAN pools in the virtual AP profile for large networks that require more than one subnet for HMD clients within a specific floor or building.

1. This virtual AP parameter enables the AP to use either 1 vlan or multiple vlans (vlan pooling) for its SSID.	<pre>! wlan virtual-ap "&lt;HMD client virtual ap name&gt;"     vlan &lt;HMD vlan # or list of vlans&gt;<sup>1</sup> !</pre>
---	--

### Fast-Roaming Technologies (OKC)

- Use WPA2-AES Enterprise as the opmode for the SSID Profile to provide the best security and roaming enhancement feature for HMD clients.

1. This SSID profile parameter enables the AP to use the strongest encryption type for wireless with its SSID which supports opportunistic key caching (OKC).	<pre>! Wlan ssid-profile "&lt;HMD client ssid profile name&gt;"     essid "&lt;HMD client ssid name&gt;"     opmode wpa2-aes<sup>1</sup> !</pre>
2. This dot1x profile parameter enables the Aruba Mobility Controller to use OKC for WPA2-AES enterprise clients that support it so that they don't have to do a full 802.1X EAP exchange when roaming across APs on the same Mobility Controller.  <b>NOTE:</b> OKC does not work across controllers.	<pre>! aaa authentication dot1x &lt;HMD dot1x profile name&gt;     opp-key-caching<sup>2</sup> !</pre>

- For clients that don't support OKC this dot1x profile parameter informs the Mobility Controller to look at the 802.11 Association Request to see if a pmkid is present. If the value is 0 then there will be a full 802.1X EAP exchange. For example, this feature should improve dot1x roaming with Apple Macbooks that associate to WPA2-AES SSIDs when OKC is enabled in the dot1x profile.

```
!
validate-pmkid3
!
```

## IP Mobility Configuration

Layer 2 or layer 3 mobility configuration is addressed in Aruba base designs, and is typically considered prior to optimizing for HMDs. Some Aruba networks begin with L2 before having any HMDs, and later find that L3 is now required due to roaming patterns of new devices. That is the context in which IP Mobility is addressed in this Guide.

To plan a mobility domain, begin by taking a look at the network map, with a special focus on the APs and controllers. Generally, this will provide the information you need to develop a logical grouping of mobility domains. You should also examine heat maps of your network, and determine if the coverage areas provide enough connectivity and overlap to allow your clients to transition networks. Outdoor APs may extend this coverage between buildings, providing you with a larger mobility domain.

The L2 or L3 mobility selection is made in the Virtual AP Profile.

If L3 mobility has been selected, remember to create HAT table entry for each subnet that is in use. This is a Global CLI level command. If L3 mobility has been selected then please review the design considerations in [Layer 3 \(IP\) Mobility on page 36](#).

## IP Multicast Optimization

In this section we show how to implement the design principles related to IP Multicast Optimization.

- Enable IGMP snooping on every HMD client VLAN so that multicast traffic is only replicated to APs with active members of a multicast group, which limits the unnecessary flooding to all APs on a given mobility controller.
- Enable multicast rate optimization ("mcast-rate-opt") so that every broadcast and multicast frame that is sent in the air to HMD clients is sent at the lowest associated device 802.11 data rate, which is generally higher than the lowest configured data rate for the specified PHY type.

- Apply IGMP snooping to all HMD client VLANs on all mobility controllers to make sure only necessary traffic is sent to the air.
- This SSID profile parameter forwards broadcast and multicast packets in the air at the highest 802.11 control data rate. Apply this feature to all HMD SSID profiles to provide the best performance to associated clients.

```
!
interface vlan <vlan number for every active
vlan>1
    ip igmp snooping
!
wlan ssid-profile "<HMD client ssid profile name>"
    mcast-rate-opt2
!
```

## Interference Resistance

The next few pages detail how to implement the design principles related to Interference Resistance for all HMD classes.

### Physical and RF Design Optimizations

If interference is suspected in your environment, begin by working through the checklist of physical environment changes that were detailed in Chapter 3:

- Verify that the AP placement meets Aruba and industry standard best practice guidelines, and is neither too sparse nor too dense.
- Make sure that FHSS APs are located at least 20 feet from all APs.
- Ensure that the uniform SNR in the environment exceeds the minimum requirement of the HMD manufacturer.
- Reconfigure the output power and/or channel of fixed-frequency interferers.
- Relocate, phase out, or otherwise eliminate identified sources of interference.



In the case of many video and voice devices, this task can be accomplished by converting the devices to use 802.11 instead of proprietary radio schemes.

The balance of this section addresses controller optimizations that enhance the system's ability to perform in the presence of elevated noise levels.

### Mode-Aware ARM

- Enable co-channel interference (CCI) mitigation in the RF ARM profile so that the APs can find the best coverage pattern in each RF neighborhood.

1. This ARM profile feature allows the APs to dynamically choose the best coverage pattern of APs and air monitors.	<pre>! rf arm-profile "&lt;arm profile name&gt;"   mode-aware<sup>1</sup> !</pre>
---	---

In addition to enabling these features, Aruba recommends limiting the minimum and maximum transmit power settings that ARM can use. It is important to match client and AP power. This is especially important for voice devices that typically have limited batteries and radio output.

1. In the ARM profile that is assigned to the radio profile that the voice HMD clients will associate to it is suggested to configure the minimum transmit value to be at least 18 dBm.	<pre>! rf arm-profile "&lt;HMD client arm profile name&gt;"   min-tx-power &lt;desired minimum transmit power&gt;<sup>1</sup></pre>
---	---

## Basic Rates

- If the handset supports only 802.11b, set the basic rates to 1 and 2, and set the supported rates on the APs to 1, 2, 5.5, and 11. Lower basic rates increase reliability in certain cases where the client may have issues receiving acknowledgements at higher rates in a dense environment, or if the client is in the cell border.

1. This SSID profile configuration informs all clients who see this SSID to connect only at 802.11b rates.	<pre>! Wlan ssid-profile "&lt;HMD client ssid profile name&gt;"   g-basic-rates 1 2<sup>1</sup>   g-tx-rates 1 2 5 11<sup>1</sup> !</pre>
--	---

## Max-Retries

A general best practice for high interference environments is to set the retries on the controller and handset to 15.

1. This SSID profile suggested configuration as stated above helps improve the connection performance to the client.	<pre>! Wlan ssid-profile "&lt;HMD client ssid profile name&gt;"   max-retries 8<sup>1</sup> !</pre>
--	---

## Noise Immunity

As explained in Chapter 3, a new feature called Noise Immunity has been introduced in the rf dot11g-radio-profile in ArubaOS 3.3.2.11. This feature works only in the 2.4 GHz band.

- Level-0: No ANI Adaptation
- Level-1: Noise Immunity only
- Level-2: Default setting (only Spur and Noise Immunity)
- Level-3: Disables weak OFDM detection (if Spur and Noise Immunity fails to mitigate interference)
- Level-4: Disables First Step (FIR) (if Spur, Noise Immunity and disabling weak OFDM detection fails to mitigate interference)
- Level-5: Disables PHY error reporting in cases of high interference if all the above steps fail to reduce interference

1. Example Noise Immunity setting that raises the default value to disable weak OFDM detection and FIR.	<pre>! rf dot11g-radio-profile &lt;802.11g rf profile name&gt;   interference-immunity Level-4 !</pre>
---	--

This chapter explains the handful of special optimizations for single-purpose HMDs (such as printers, scanners, patient monitors, and so on) to operate most effectively with Aruba controllers and APs.



All Multi-Purpose device configuration recommendations described in [Chapter 4, “Configuring Global Settings for All Roaming Devices”](#) also apply to Single-Purpose devices. These settings should be applied before proceeding with this section.

Single-purpose HMDs do not require incremental configuration of the Airtime, IP Mobility, IP Multicast or Interference Resistance parameters. This chapter considers only additional changes to Device Configuration and Roaming Optimization for this class of device.



The changes specific to the single-purpose HMD (DTIM, Tx rates) should be applied to only the isolated HMD SSID (not to other SSID profiles)

## Device

### Device Configuration

If the manufacturer of the single-purpose HMD recommends a specific DTIM value, you must implement this on both the controller and the device.

1. This SSID profile parameter helps provide more battery life by transmitting less to the HMD client when it is in power save mode so that it does not have to wake up for every packet that leaves from the AP.	<pre>! wlan ssid-profile "&lt;HMD client ssid profile name&gt;"     dtim-period &lt;HMD client manufacturer suggested DTIM value&gt; !</pre>
---	--

## Roaming Optimization

This paragraph shows the configuration statements required to implement the design principles related to Roaming Optimization for single-purpose HMDs.

1. This setting helps HMD clients with very sensitive roaming algorithms that require very strong AP signal strength. Apply this setting to the ARM profile in both 802.11a and 802.11g radio profiles. By default, the same ARM profile is tied to both radio profiles.
2. To provide better interoperability with older HMD clients, it is suggested to limit the 802.11 data rates. Apply this setting to all single-purpose HMD clients.

```
!  
rf arm-profile "<HMD client arm profile name>"  
    min-tx-power 181  
!  
wlan ssid-profile "<HMD client ssid profile name>"  
    g-tx-rates 1 2 5 6 9 11 12 18 242  
!
```



This chapter explains the additional special optimizations recommended by Aruba to configure voice HMDs to operate most effectively with Aruba controllers and APs. The configuration parameters below are in addition to what is stated in the previous sections for Multi-Purpose HMDs.



All Multi-Purpose device configuration recommendations described in [Chapter 4, “Configuring Global Settings for All Roaming Devices”](#) also apply to Voice devices. These settings should be applied before proceeding with this section.

## Device

### Device Configuration

This section reviews how to implement the design principles related to Device Configuration.

#### Maximize Handset Battery Life

Power-saving mechanisms help improve the battery life on the handsets by allowing the handsets to sleep longer. This service is part of the Aruba Voice Services Module license and is strongly recommended if voice is a production application in your environment.

If the handset supports Unscheduled, Asynchronous Power Save Delivery (UAPSD), set DTIM to 10 to enable UAPSD on the handset and enable WMM on the infrastructure.

1. This SSID profile suggested configuration as stated above helps improve battery life on the handsets that support such stated features as WMM and U-APSD.	<pre>! Wlan ssid-profile "&lt;HMD client ssid profile name&gt;"   dtim-period 10<sup>1</sup>   wmm<sup>1</sup>   wmm-uapsd<sup>1</sup> !</pre>
--	--

For handsets that do not support UAPSD, set the DTIM to 3. Another option that could save power on that handset is to enable the battery boost feature which converts all broadcast/multicast traffic into unicast traffic. However, that feature should be tested during a maintenance window to make sure there are no issues found when it is enabled.

1. This SSID profile suggested configuration as stated above helps improve battery life on the handsets by converting all multicast traffic into unicast traffic. This allows the possibility of configuring a larger DTIM.	<pre>! Wlan ssid-profile "&lt;HMD client ssid profile name&gt;"   dtim-period 3<sup>1</sup>   battery-boost<sup>1</sup> !</pre>
---	---

## End-to-End QoS

Quality of service (QoS) has already been enabled on the Aruba controller with the WMM command in the SSID Profile just shown.

As discussed in the Chapter 5, all network elements in the LAN and WAN must be configured to support QoS (802.1p or DSCP) if voice traffic is present.

Finally, ensure that the voice handsets themselves are properly configured, and that the settings match the ones used on the infrastructure.

## Airtime Optimization

This section reviews how to implement the design principles related to Airtime Optimization for voice devices.

### Complete a Voice Capacity Plan

Use the capacity planning strategy in conjunction with the recommended per-AP simultaneous call values in [Table 5 on page 27](#), to ensure that there are sufficient APs in the environment to handle the offered load of voice devices.

Aruba recommends that each customer test call capacity in a lab environment with their specific mix of handsets, APs, wired infrastructure, and back-end call servers to validate the design prior to deployment.

### Enable Call Admission Control (CAC)

Aruba strongly recommends enabling CAC for production voice deployments. The maximum number of calls supported per AP that was determined in the design phase should be programmed here. CAC is implemented on a per AP, per radio basis. Set the handoff reservations and the high capacity threshold value to 20%.

<ol style="list-style-type: none"> <li>1. Enables CAC feature. This requires a Voice Services Module license.</li> <li>2. Specifies capacity percentage reserved for mobile VoIP clients (default = 20%).</li> <li>3. Specifies remaining capacity percentage that enables roaming for new clients (default = 20%).</li> </ol>	<pre>! wlan voip-cac-profile "&lt;CAC profile name&gt;"   call-admission-control<sup>1</sup>   call-handoff-reservation &lt;percentage&gt;<sup>2</sup>   high-capacity-threshold &lt;percentage&gt;<sup>3</sup> !</pre>
--	---

## Roaming Optimization

This section reviews how to implement the design principles related to Roaming Optimization for voice devices.

### Enable ARM with Voice-Aware and Min/Max Output Power

As explained in Chapter 5, ARM should be enabled for channel and power management for voice deployments.

Once ARM is running, enable voice aware scanning. In addition to enabling these features, Aruba recommends limiting the minimum and maximum transmit power settings that ARM can use. Consult [Table 7 on page 32](#) for specific minimum/maximum bands recommended by Aruba engineers.

1. In the ARM profile that is assigned to the radio profile that the voice HMD clients will associate to, it is suggested to enable voip-aware-scan and configure the maximum transmit value to match the voice HMD maximum transmit power.  If there will be Voice over Wi-Fi HMD clients, it is suggested to enable ARM profile feature “voip-aware-scan” so that phones on calls will not be disrupted with APs periodically scanning other 802.11 channels until the call is finished. This feature requires the Voice license.	<pre> ! rf arm-profile "&lt;HMD client arm profile name&gt;"   voip-aware-scan<sup>1</sup>   min-tx-power &lt;desired minimum transmit power&gt;<sup>1</sup>   max-tx-power &lt;desired maximum transmit power&gt;<sup>1</sup> </pre>
---	---

### Configure Max Retries, Max Transmit Failures, and Disable Probe Retries

#### Max-Retries

A general best practice for voice deployments is to set the retries on the controller and handset to 2. Because VoIP is delay sensitive, after the packet is delayed, retrying in order to successfully transmit a packet may just add to the latency in the network.



In noisy environments (from the radio point of view), it is recommended to increase the retries value.

#### Max Transmit Failures

Set the max-tx-fail retries value to 25. This value is the number of consecutive transmitted frames from the AP that are not acknowledged by the HMD client that the frames are destined to. If the station ACKs any frame, then the counter is reset.

## Disable Probe Retries

In the current ArubaOS, 802.11 Probe Response retries keep the same timestamp as the first Probe Response. This can cause clients to get out of sync with the AP, so it is a best practice to disable probe retries. Please note this feature requires a voice license.

1. This SSID profile suggested configuration as stated above helps improve the connection performance to the client.	!
2. This parameter will not be available until the Voice license is installed.	Wlan ssid-profile "<HMD client ssid profile name>" max-retries 2 <sup>1</sup> max-tx-fail 25 <sup>1</sup> disable-probe-retry <sup>2</sup> !

## Polycom SpectraLink 8020/8030 Wireless Telephones Configuration Template

The following CLI configuration output has been tested in several Aruba/Polycom customer sites. It is also recommended to follow the *Polycom Best Practice Guide to Deploying SpectraLink 8020/8030 Wireless Telephones* that is located on their website ([http://www.polycom.com/support/voice/wi-fi/spectralink\\_8030\\_wireless.html](http://www.polycom.com/support/voice/wi-fi/spectralink_8030_wireless.html)). For example, most of the customer sites with the configuration shown below had a deployment coverage with at least -67 dBm signal strength for the voice SSID, and the APs were configured to broadcast the voice SSID at 802.11a only with a configured maximum transmit power of 18 dBm. Polycom recommends that the SVP Server and the voice handsets be deployed on the same subnet, so make sure that they are both on the same VLAN or verify that multicast traffic is routable to the server.

```
!
wlan virtual-ap "<SpectraLink VAP name>"
    allowed-band a
    ssid-profile "<SpectraLink SSID profile name>"
    vlan <vlan #>
    aaa-profile "<SpectraLink aaa profile name>"
!
aaa derivation-rules user <user derivation rule name>
    set role condition essid equals "<voice SSID name>" set-value voice
!
aaa profile "<SpectraLink aaa profile name>"
    initial-role "logon"
    user-derivation-rules <user derivation rule name>
!
wlan ssid-profile "<SpectraLink SSID profile name>"
    essid "<SpectraLink SSID name>"
    opmode wpa-psk-tkip
    dtim-period 2
    a-basic-rates 6 12 24
    a-tx-rates 6 9 12 18 24
    max-retries 3
    wpa-hexkey <key value>
    max-tx-fail 25
    disable-probe-retry
    no wmm
!
rf dot11a-radio-profile "<802.11a radio profile name>"
    arm-profile "<802.11a arm profile name>"
!
rf arm-profile "802.11a arm profile name"
    max-tx-power 18
    min-tx-power 12
    voip-aware-scan
!
```

```

user-role voice
  session-acl sip-acl
  session-acl noe-acl
  session-acl svp-acl
  session-acl vocera-acl
  session-acl skinny-acl
  session-acl h323-acl
  session-acl dhcp-acl
  session-acl tftp-acl
  session-acl dns-acl
  session-acl icmp-acl
!

```

## Cisco 792x Series Phones Configuration Template

The following CLI configuration output has been tested in several Aruba/Cisco 7921g and 7925 customer sites. It is also recommended to follow the Cisco deployment guide that is located on their website ([http://www.cisco.com/en/US/docs/voice\\_ip\\_comm/cuipph/7921g/6\\_0/english/deployment/guide/7921dply.pdf](http://www.cisco.com/en/US/docs/voice_ip_comm/cuipph/7921g/6_0/english/deployment/guide/7921dply.pdf)). For example, most of the customer sites with the configuration shown below had a deployment coverage with at least -67 dBm signal strength for the voice SSID, and the APs were configured to broadcast the voice SSID at 802.11a only with a configured maximum transmit power of 18 dBm.

```

!
wlan virtual-ap "<792X VAP name>"
  allowed-band a
  ssid-profile "<792X SSID profile name>"
  vlan <vlan #>
  aaa-profile "<792X aaa profile name>"
!
aaa derivation-rules user <user derivation rule name>
  set role condition essid equals "<voice SSID name>" set-value voice
!
aaa profile "<792X aaa profile name>"
  initial-role "logon"
  user-derivation-rules <user derivation rule name>
!
aaa authentication dot1x "<dot1x authentication profile>"
  no opp-key-caching
!
wlan ssid-profile "<792X SSID profile name>"
  essid "<792X SSID name>"
  opmode wpa2-aes
  dtim-period 2
  a-basic-rates 6 12 24
  a-tx-rates 6 9 12 18 24
  max-retries 3
  max-tx-fail 25
  disable-probe-retry
  wmm
  deny-bcast
!
rf dot11a-radio-profile "<802.11a radio profile name>"
  arm-profile "<802.11a arm profile name>"
!

```

```


rf arm-profile "802.11a arm profile name>"
  max-tx-power 18
  min-tx-power 12
  voip-aware-scan
!
rf ht-radio-profile <802.11a ht radio profile>
  single-chain-legacy
!
user-role voice
  session-acl sip-acl
  session-acl noe-acl
  session-acl svp-acl
  session-acl vocera-acl
  session-acl skinny-acl
  session-acl h323-acl
  session-acl dhcp-acl
  session-acl tftp-acl
  session-acl dns-acl
  session-acl icmp-acl
!

```

## Vocera

The Vocera solution shown in [Figure 13](#) consists of wearable Vocera Communications Badges with integrated 802.11b radios in conjunction with the Vocera Communications software server.

**Figure 13** Vocera Integration Into Aruba Network



arun\_130



VoIP call management and speech recognition engine functionality are incorporated in the Vocera Communications server software, which runs on standard Windows servers. Through the use of optional modules, Vocera can interface with circuit PBX, alarm/alert, and nurse call systems. Medical staff can log into the system using voice commands and can call colleagues by name or role (for example, radiologists or charge nurses) by speaking into the badge. The system supports WPA (802.1X authentication and TKIP encryption) for secure voice communication. The Aruba user-centric network supports fast roaming (<100 ms) of Vocera badges, which significantly improves call quality and reduces the occurrence of “dropped calls” due to latency. The user-centric network, with its unique awareness of the application layer, is able to recognize badges from their use of the Vocera VoIP protocol. This capability allows Aruba to ensure optimal load through call admission control of the voice badges. Furthermore, the user-centric network supports Voice-Aware Scanning, which is the ability to postpone AP scanning for RF management/security purposes in the presence of a Vocera badge to ensure that QoS is not negatively impacted. The Vocera system is supported on any Aruba mobility controller and any AP with an 802.11b/g radio.

## Vocera Configuration Template

**Table 8** provides a summary of required multicast settings and best practices for Vocera system implementations.

**Table 8** *Multicast Recommendations for Vocera System Implementations*

Multicast Settings	
Not Using IGMP	• Must be 12 only network
	• IGMP snooping must be disabled on all L2 devices in the audio path
	• Enable IGMP on Badge
IGMP	• IGMP V2
	• Multicast routing must be enabled on all routers between WLC and APs
	• PIM (sparse mode or sparse-dense-mode) must be enabled on all router interfaces between WLC and APs
	• IGMP snooping should be enabled on all L2 devices
	• Enable IGMP on Badge

The following CLI configuration output has been tested in several Aruba/Vocera customer sites. It is also recommended to follow the Vocera best practice guides (Badge Configuration Guide and Infrastructure Planning Guide) that are located on their website (<http://www.vocera.com/documentation/default.aspx>). For example, most of the customer sites with the configuration shown below had a deployment with at least -65 dBm signal strength for the voice SSID, and the APs were configured with a 3 channel plan on 802.11bg radios with a configured maximum transmit power of 15 dBm.

```
!
wlan virtual-ap "<Vocera VAP name>"
  allowed-band g
  ssid-profile "<Vocera SSID profile name>"
  vlan <vlan #>
  aaa-profile "<Vocera aaa profile name>"
!
aaa derivation-rules user <user derivation rule name>
  set role condition essid equals "<Vocera SSID name>" set-value voice
!
```

```

!
aaa profile "<SpectraLink aaa profile name>"
    initial-role "logon"
    user-derivation-rules <user derivation rule name>
!
wlan ssid-profile "<Vocera SSID profile name>"
    essid "<Vocera SSID name>"
    mcast-rate-opt
    opmode wpa-psk-tkip
    dtim-period 1
    g-basic-rates 1 2
    g-tx-rates 1 2 5 6 9 11 12 18 24
    max-retries 4
    wpa-hexkey <key value>
    max-tx-fail 25
    disable-probe-retry
    wmm
!
rf dot11g-radio-profile "<802.11g radio profile name>"
    arm-profile "<802.11g arm profile name>"
!
rf arm-profile "802.11g arm profile name"
    max-tx-power 15
    min-tx-power 12
    voip-aware-scan
!
ip access-list session vocera-acl
    any any svc-vocera permit queue high
    any any udp 5001 permit queue high
    any any udp 5005 permit queue high
    any any udp 5100 5200 permit queue high
    any any udp 5251 permit queue high
    any any udp 5300 5400 permit queue high
    any any udp 5555 5556 permit queue high
    any any tcp 5555 5556 permit queue high
!
user-role voice
    session-acl sip-acl
    session-acl noe-acl
    session-acl svp-acl
    session-acl vocera-acl
    session-acl skinny-acl
    session-acl h323-acl
    session-acl dhcp-acl
    session-acl tftp-acl
    session-acl dns-acl
    session-acl icmp-acl
!

```

## Ascom i75 Phones Configuration Template

The following CLI configuration output has been tested in several Aruba/Ascom customer sites. It is also recommended to follow Ascom's best practice guide that is located on their website <http://www.ascomwireless.com/pdf/guide/freeNETvowifiSystemPlanningGuide-vPA.pdf>. Most of the customer sites with the configuration below had a deployment coverage with at least -67 dBm signal strength for the voice SSID and the APs were configured to broadcast the voice SSID at 802.11g only with a configured maximum transmit power of 15 dBm across a 3 channel plan deployment.

```
!
wlan virtual-ap "<Ascom Virtual AP name>"
    allowed-band g
    ssid-profile "<Ascom SSID profile name>"
    vlan <vlan #>
    aaa-profile "<Ascom aaa profile name>"
    broadcast-filter all
    broadcast-filter arp

!
aaa derivation-rules user <user derivation rule name>
    set role condition essid equals "<voice SSID name>" set-value voice
!
aaa profile "<Ascom aaa profile name>"
    initial-role "logon"
    user-derivation-rules <user derivation rule name>
!
wlan ssid-profile "<Ascom SSID profile name>"
    essid "<Ascom SSID name>"
    opmode wpa2-psk-aes
    dtim-period 5
    a-basic-rates 6 12
    a-tx-rates 6 9 12 18 24
    max-retries 4
    wpa-passphrase "<passphrase value>"
    max-tx-fail 25
    disable-probe-retry
    wmm

!
rf dot11g-radio-profile "<802.11g radio profile name>"
    arm-profile "<802.11g arm profile name>"
    802.11h

!
rf arm-profile "<802.11g arm profile name>"
    max-tx-power 15
    min-tx-power 3
    voip-aware-scan

!
user-role voice
    session-acl sip-acl
    session-acl noe-acl
    session-acl svp-acl
    session-acl vocera-acl
    session-acl skinny-acl
    session-acl h323-acl
    session-acl dhcp-acl
    session-acl tftp-acl
    session-acl dns-acl
    session-acl icmp-acl
!
```

Troubleshooting client device issues in wireless networks consists of systematically narrowing down the source of the problem by knowing all the pieces involved in providing seamless mobility to HMDs. This chapter documents the processes used by senior Aruba support engineers to resolve problems with roaming devices. It will help you to identify and troubleshoot the most common problems found in WLAN connectivity.

## Scoping the Problem

The first step is to have a clear understanding of the issue being reported so that the next steps can be efficiently chosen. [Table 9](#) lists several symptoms and possible causes to help you initially scope the problem.

**Table 9** *Trouble Symptoms and Causes*

Symptom	Possible Cause
Issue is isolated to an individual	Might be a NIC, supplicant, or driver related issue
Issue is isolated to a geographical area	Might be a RF or other physical layer problem
Issue affects a group of people on a certain SSID	Might be an AP configuration problem
Issue affects a group of people on a common group of APs	Might be an AP configuration or L2/L3 problem
Issue is isolated to a certain application	Might be a routing problem or an application layer problem
Issue is isolated to a particular server	Might be a routing or server problem
Issue is isolated to a particular time of day	Might be a non-802.11 device, firewall, or service issue

Other things to check:

- Has anything changed in the WLAN equipment configuration? (All the Aruba Mobility Controllers have an audit log that tracks every GUI and CLI configuration change.)
- Has anything changed in the network?
- Has anything changed in the area of the reported problem?

## Mobility Framework

As we have seen, designing a WLAN for highly mobile devices involves many hardware and software components, all operating in the most optimal manner. Troubleshooting these highly active WLANs requires skills learned by understanding the sequence of protocols involved in providing end-to-end connectivity and by having the experience in checking common symptoms to complete the process of elimination. It is therefore imperative to know what areas can affect wireless mobility. See [Table 10](#) for a basic list of network elements and their corresponding components.

**Table 10** *Possible Component Trouble Spots*

Component	Things to Check
Wireless Client	<ul style="list-style-type: none"><li>• Device hardware</li><li>• Device OS</li><li>• Device supplicant</li><li>• Device driver</li></ul>
Access Point (AP)	<ul style="list-style-type: none"><li>• AP physical location</li><li>• Antenna position</li><li>• AP status</li><li>• AP configuration</li></ul>
Backend Servers	<ul style="list-style-type: none"><li>• DHCP server</li><li>• RADIUS server</li><li>• LDAP server</li><li>• User database (for example, Microsoft Active Directory)</li></ul>

## HMD Troubleshooting

Upon receiving a report of an HMD connectivity issue, be sure to gather the following information:

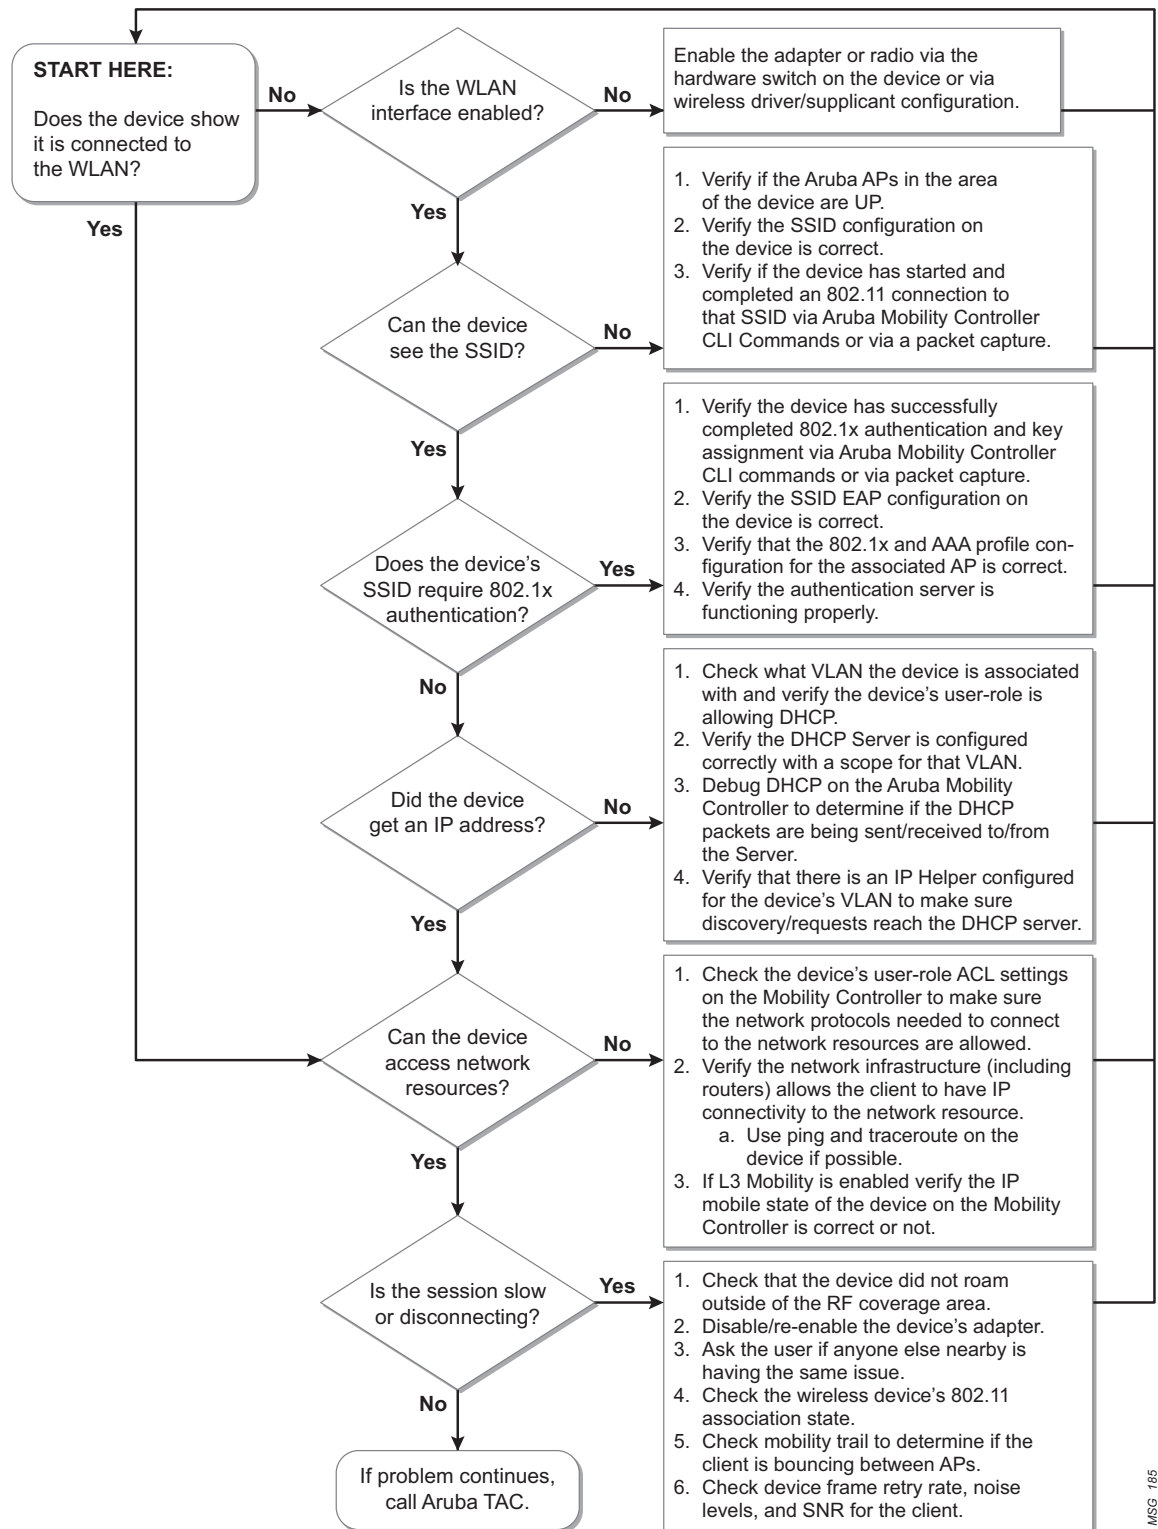
1. Device location (country, city, building, floor, general location, office number)
2. Device username (if using L2 or L3 authentication)
3. Device NIC MAC address
4. Device IP address (if available)

The location determines the Aruba Mobility Controller(s) on which to concentrate troubleshooting efforts.

## Troubleshooting Flow Chart

The flow chart shown in [Figure 14](#) shows you where to start troubleshooting when a problem is reported for a mobile device. The sections following the flowchart help explain in further detail what steps can be taken to help resolve or narrow down the issue.

**Figure 14** *Roaming Device Troubleshooting Flowchart*



MSG\_185

## Symptom #1—Device cannot see any SSIDs

Suggested Action:

- Check building location of the wireless device.
- SSH to the Aruba Mobility Controller responsible for the building.
- Verify that APs are up on the Controller.
  - **show ap bss-table** and **show ap active**  
This command lists all APs with their respective AP names and their active SSIDs and BSSIDs.
- Verify that the SSID is not hidden.
  - If the SSID is hidden, verify that the client is properly configured to associate to it.
- Check wireless NIC enable/disable physical switch on mobile device.
- Check wireless NIC enable/disable soft setting within wireless supplicant software.
- If the device is using Microsoft Windows Operating System, issue a “repair” in Network Connections or wireless NIC system tray icon.

If no issues are found and the above actions have not corrected the problem, continue with device troubleshooting. In addition to previously gathered information (username, location, MAC/IP), gather device hardware model name/number and wireless NIC brand/model/type/driver version for further troubleshooting. Also, take a wireless packet capture so that Aruba Support can perform analysis by means of the AP Remote Packet Capture method or by means of third-party software (for example, Wireshark, Wireshark, Wireshark, and so on). Please also provide the Aruba Support Team all the necessary CLI command output for Mobility Controller, AP, and user statistics.

## Symptom #2—Device can see some SSIDs but not the one to which it needs to connect

Suggested Action:

- Verify that the required SSIDs are active and enabled in the Aruba Mobility Controller.
  - **show ap bss-table**  
Issue this command from any mobility controller, master or local that is servicing APs. This command lists all APs with their AP names and their active SSIDs and BSSIDs.
- Verify that all APs are up and active, especially those in the area of the problem device.
  - **show ap database long**  
Issue this command from the master mobility controller servicing the area of the problem device. This command lists all known APs serviced from that master mobility controller, regardless of being up or down.
- If all APs have proper SSID configurations and no APs are reported down, verify that the client device is attempting to associate/authenticate.
  - **Configure terminal**
    - **logging level debugging user-debug <wireless device's mac address>**
    - **end**

This command starts debugging on all Aruba processes for the wireless device and logs the results in Aruba logging category “user-debug”

View the debug output with the CLI command “**show log user-debug all | include <wireless device's mac address>**”



- **show ap debug mgmt-frames client-mac <wireless device's mac address>**

This lists the 802.11 management packets (Association Request, Association Response, Re-Association Request, Re-Association Response, Disassociation, and Deauth) for the specified wireless device.

If you see the latest packet as “assoc-resp,” the wireless device should be authenticating (if VAP is configured for L2 or L3 authentication) or should be authenticated already.

- **show log system all | include "<wireless device's mac address>"**

Issue this command from any master or local mobility controller that is servicing APs to which the device may attempt to associate. This command shows if the problem client is attempting to associate. Look for the problem client MAC address. It also shows to which AP the client MAC is attempting to associate. Note the BSSID.

- **show ap association client-mac "<wireless device's mac address>"**

This command shows the 802.11 state of the wireless device, what SSID it is associated to, what VLAN it is assigned, what PHY type it is using, how long it has been associated to the AP's BSSID, and what capabilities it has such as WMM, Active/Not Active, RRM client, Band Steerable, or HT capable.

Use the AP BSSID and device MAC taken from this command.

- **show ap association | include "<AP BSSID that the device is associated to>" and "show user-table bssid "<AP BSSID that the device is associated to>"**

This output can be used to verify if there are other devices currently associated to the same AP, thus helping to rule out infrastructure issues as compared to a single-client issue.

- **show log security all | include "<wireless device's mac address>"**

Look for the problem client MAC. This command can be used to determine whether the client is attempting to authenticate via L2 or L3 authentication and if the request is being rejected. If the attempt is rejected, this can be established as the reason for client failure. Investigate authentication server logs as needed.

- **show auth-tracebuf**

If the device is configured to use L2 authentication such as 802.1X, verify the wireless device successfully completed all EAP and Key exchange phases with the CLI command

**"show auth-tracebuf mac <wireless device mac address>"**

- **show log errorlog all**

This command can be used to determine if there are any miscellaneous errors with the mobility controller, the AP, or the wireless device.

This command can also point to problems with an authentication server not responding to authentication requests if L2/L3 authentication is enabled on the virtual AP to which the device is trying to connect.

If the authentication server is RADIUS, look for excessive RADIUS timeouts or instances of the Aruba Mobility Controller taking a RADIUS server out of service for the server hold-down timer (this indicates possible RADIUS server connectivity or performance issues and should be investigated as needed).

Using these steps, you can determine if the device has passed 802.11 negotiation and is attempting to authenticate (if L2/L3 authentication is required). If none of the above steps yields correcting information, take a wireless packet capture for Aruba Support to analyze. You can use the AP Remote Packet Capture method or third-party software (for example, Wildpackets Omnipcap, Cace Technologies Airpcap, and so on). Please also provide the Aruba Support Team all the necessary CLI command output for Mobility Controller, AP, and user statistics.

## Symptom #3—Device successfully authenticates but cannot communicate

This scenario is most likely related to the device being in a restricted user-role (firewall ACL misconfiguration) or it is not getting an IP address from the DHCP server due to VAP configuration, DHCP connectivity issues, DHCP Scope misconfiguration, or L3 Mobility issues (if enabled).

Suggested Action:

- Verify that the device is receiving an IP address via device statistics or via the Aruba Mobility Controller:
  - **show user mac <wireless device MAC address>**

This command displays all details pertaining to the client. Verify that the IP address is not 0.0.0.0 or a 169.x.x.x address.

This command is also used to verify if the user was successfully authenticated, and displays the user-role, ACL number, authentication method, and associated AP name/BSSID.

If the device is associated to the right user-role and VLAN but does not have a valid IP address, disable/re-enable its wireless adapter or force a DHCP 'release-renew' in the device's operating system.

If the problem is not corrected, investigate DHCP infrastructure and connectivity.
  - DHCP Troubleshooting:

Enable DHCP debugging on the Aruba Mobility Controller at the AP device location.

    - **Config t**
    - **logging level debugging network subcat dhcp**
    - **end**

View the DHCP debug for the wireless device using the CLI command

```
"show log network all | include <wireless device MAC address>
```

Confirm that the DHCP server is in service.

Verify that the upstream router has the correct DHCP helper-address for the device's VLAN.

Investigate whether or not the DHCP scope is correctly configured for the device's subnet and that it has available IP addresses in its pool.
- Verify that the device has been placed in the correct user role with the correct session policies.
  - **'show user mac <wireless device MAC address>' or 'show user ip <ipaddr>'**

This command lists all details pertaining to the client. Use this output to confirm that the user's authenticated role is correct.
  - **show rights <device's assigned role name>**

Use this command to confirm which policies are associated to the device's authenticated role and verify that they allow the required protocols for device IP and application connectivity.
  - **show datapath session table <device IP address>'**

This command displays all IP flows between the device and the network.

Have the device attempt a connection to its required network resource and use this command to confirm traffic passing from the device is not being denied by the Aruba stateful firewall role-based policies by verifying no IP flow is marked with the "D" flag (denied).

Using these steps, you can determine if the device has received a proper IP address, has been placed in the correct user-role with the correct policies, and verify network connectivity. If none of the above steps yields correcting information, then prepare a wired packet capture for the Aruba Support team to analyze between the Aruba Mobility Controller and the uplink switch. This can be done with built-in operating system applications like tcpdump, network monitor, or third-party software like Wireshark,

Ethereal, or Wireshark's OmniPeek/EtherPeek. Another method to achieve device packet capture is by implementing session mirroring in the device's user-role on the mobility controller.

## Symptom #4—Device successfully authenticates and can communicate but is experiencing dropped connections and/or poor performance

Suggested Action:

- Confirm with the user that they did not roam outside of the engineered RF coverage area with their device.
- Disable/re-enable the device's adapter and verify if the issue persists.
- Confirm that the AP to which the device is associated is nearby.

Use the CLI command **"show ap debug client-table ap-name <ap name that the device is associated to>"** to determine the "last Rx SNR" value of the device.

Anything with "Last Rx SNR" value of 25 or greater normally provides good performance with the higher supported 802.11 data rates.

- Compare the problem user's stated location with the building and AP floor plan or use Aruba RF Plan.
- Ask the user reporting the trouble if anyone else nearby is having the same issue. This information assists in determining if this is an infrastructure or single-user problem.
- Check the user log and the AP 802.11 management frames for possible cause of disconnection.

- **show log user all ' | include "<wireless device's mac address>"**
- **show ap debug mgmt-frames client-mac <wireless device's mac address>**

This command determines from when and where the disconnection originated (either the AP or the device) and helps determine the reason.

- Check the wireless device's 802.11 association state
- **Show ap debug client-table ap-name <Aruba AP name where the wireless device is associated to>**

Part of this CLI output displays the wireless device's Last\_Rx\_SNR, Tx\_Rate, and Rx\_Rate.

If the SNR is 15 or lower, the wireless device is possibly too far from the AP. This might be due to the device's roaming algorithm not being optimal and needs to be forced to look for a closer AP by disabling/re-enabling its network adapter.

If the Tx\_Rate or Rx\_Rate are 1, 2, or 6, the device may be experiencing interference or is too far away from the AP.

If the Tx Retry rate is constantly 35% or higher, the device may be experiencing interference or is far away from the AP.

There might be non-802.11 interference if the MAC and PHY errors are at an aggregate of 20% or higher, which can be seen through the CLI command **"show ap arm rf-summary ap-name <Aruba AP name where the wireless device is associated to>"**

- Check mobility trail to determine if the client is bouncing between APs even when stationary.

- **show ip mobile trail <wireless device MAC address>**

"router mobile" must be configured in order for this CLI command to work.

This command displays the mobility history of a given client. This can be used to check for the frequency of roaming.

- Check device frame retry rate, noise levels, and SNR for the client.

- `'show user mac <device wireless MAC address>' or 'show user ip <device IP address>'`

Investigate the following:

- Channel Frame Retry Rate:
  - 10-20% is normal
  - 20-30 is intermediate
  - 40+% is very high

This means 40% of the frames sent to the air have been retransmitted.

This is a symptom of heavy interference or low signal strength between the device and the AP.

Take a wireless packet capture to see if the 802.11 frame retries are due to the AP not hearing the wireless device, or the wireless device is not hearing the AP due to interference, or the device is too far from the AP.

- Channel Noise:

If channel noise is at a value of 75 or below, this is a critical interference level that should be viewed with a Spectrum Analyzer.

From these steps you can determine possible causes for poor performance or roaming issues due to device driver sub-optimal performance, roaming outside of the WLAN coverage area, or interference. If none of these steps yields correcting information, then take a wireless packet capture for Aruba Support to analyze by means of the AP Remote Packet Capture method or third-party software (for example, Wireshark, Wireshark, Cace Technologies Airpcap, and so on). Please also provide the Aruba Support Team with all the necessary CLI command output for mobility controller, AP, and user statistics.

## Before you contact Aruba Support

In order for Aruba Support to provide the fastest problem resolution to any HMD connectivity or performance issue, the following information should be provided.

1. Provide the Aruba WLAN Controller logs and output of `show tech-support`.

CLI Example:

- a. `tar logs tech-support`
- b. `copy flash: logs.tar tftp:<tftp server IP address> <file name>`

2. Provide the Syslog Server file of the Aruba WLAN Controller at the time of the problem.

If no Syslog Server is available to capture log output from the Aruba WLAN Controller, please set one up as soon as possible, as this is a strongly suggested troubleshooting and monitoring best practice.

A free Syslog server can be found at Kiwi Enterprises (<http://www.kiwisyslog.com/>).

3. State the scope of the problem as mentioned earlier in this section.
4. If there was a configuration change, please list the exact configuration steps and commands used.
5. State the date and time (if possible) when the problem first occurred.
6. Is the problem reproducible?

If the problem is reproducible, please list the exact steps taken to recreate the problem.

7. Provide the wireless device's make, model number, and its OS version, including any service packs or patches.

8. Provide the Wireless LAN Card's make, model number, driver date, driver version, and configuration on the wireless device.
9. Provide a detailed network topology:
  - a. Include all the devices in the network between the user and the Aruba WLAN Controller with IP addresses and Interface numbers, if possible.
  - b. The diagram can be formatted as Visio, PowerPoint, JPEG, TIF, etc., or it can even be hand written and then faxed to the Aruba Support Team (1-408-227-4550).
10. Provide any wired or wireless sniffer traces taken during the time of the problem.
11. Provide the following HMD statistic output on the mobility controller:
  - a. `show aaa state user <wireless client ip address>`
  - b. `show ap association client-mac <wireless device's mac address>`
  - c. `show ap debug mgmt-frames client-mac <wireless device's mac address>`
  - d. `show ap debug client-stats <wireless device's mac address> advanced`  
Run this command at least 3 times during the debugging.
  - e. `show ap monitor stats ap-name <ap name> mac <client mac> verbose`  
Run this command at least 3 times during the debugging.
  - f. `show auth-tracebuf mac <wireless client mac address>`
12. Provide the following AP statistics on the mobility controller output:
  - a. `show ap tech-support ap-name <Aruba AP name where the wireless device is associated to>`  
Run this command at least 3 times for every AP the wireless device has a problem with performance or roaming to.
13. If Layer 3 Mobility is enabled on the mobility controllers, provide the following CLI output:
  - a. `show ip mobile binding | begin <wireless device's mac address>`
  - b. `show ip mobile domain`
  - c. `show ip mobile global`
  - d. `show ip mobile host <wireless device's mac address>`
  - e. `show ip mobile remote <wireless device's mac address>]`
  - f. `show ip mobile trace <wireless device's mac address>`
  - g. `show ip mobile traffic foreign-agent`
  - h. `show ip mobile traffic home-agent`
  - i. `show ip mobile traffic proxy`
  - j. `show ip mobile traffic proxy-dhcp`
  - k. `show ip mobile trail <wireless device's mac address>`
  - l. `show ip mobile visitor <wireless device's mac address>`



To validate the Aruba device-agnostic architecture, the Aruba solution is tested with a broad set of mobile devices for interoperability, security and performance. The list of tested HMDs is updated periodically. The most recent version may be found at <http://www.arubanetworks.com/support/interoperability.php>.

## Commonly Deployed Single-Purpose HMDs

Table 11 lists some commonly deployed single-purpose HMDs that have completed interoperability testing at Aruba Networks, while Table 12 gives links to the datasheets for each device so you can gather more detailed information on each one.

**Table 11** *Single-Purpose HMDs Tested on Aruba Infrastructure*

Vendor	Device Type	Device Model	Operating System	Software Version
Symbol	MC3000	MC3090	Win CE	5.00.1400
Symbol	MC50	MC5040	Win Mobile 2003	4.21.1088
Symbol	MC70	MC7090	Win Mobile 5.0	5.1.70
Symbol	MC9000	MC9090S	Win Mobile 5.0	5.1.70
Symbol	PPT8800	PPT8846	Win CE .NET	4.10
Symbol	PPT8100	PPT8146	MS PocketPC	
Symbol	VC5090	VC5090	Win CE	5.00.1400
Symbol	MK2000	MK2046	Win CE	4.10
Symbol	WT4090	WT4090	Win CE	5.00.1400
Symbol	PDT6800	PDT6846	DOS	
Intermec	700 Series	751	MS PocketPC	4.20
Intermec	CN2	CN2B	MS PocketPC	4.20
Intermec	CN3	CN3	Win Mobile 5.0	5.1.342
Intermec	CK31	CK31	Win CE .NET	4.20
Intermec	CK60	CK60	Win Mobile 5.0	5.1.70
Intermec	T2425	T2425	DOS	
Intermec	T2455	T2455	DOS	
Intermec	CV60	CV60	Win CE .NET	4.20
Teklogix	Workabout Pro	Workabout Pro	Win CE .NET	4.20
Teklogix	7530	7530	Win CE .NET	4.20
Teklogix	7535	7535	Win CE .NET	4.20
Vocollect	Talkman T5	Talkman T5	Win CE .NET	4.20
Zebra	QL220	QL220	Embedded OS	V79.50
Zebra	RW220	RW220	Embedded OS	V90.14

**Table 12** Information Links for Certain Single-Purpose HMDs

Handheld Scanners	Datasheet Link
Symbol (Motorola) MC9090-S	<a href="http://www.motorola.com/staticfiles/Business/Products/Mobile%20Computers/Handheld%20Computers/MC9090-S/_Documents/Static%20Files/MC909x-S_1205.pdf">http://www.motorola.com/staticfiles/Business/Products/Mobile%20Computers/Handheld%20Computers/MC9090-S/_Documents/Static%20Files/MC909x-S_1205.pdf</a> <a href="http://www.leopardsystems.com.au/CartFinity/Assets/MC9090%20G_QSG.pdf">http://www.leopardsystems.com.au/CartFinity/Assets/MC9090%20G_QSG.pdf</a>
Symbol (Motorola) PDT6846	<a href="http://www.symbol.com/product.php?productID=251">http://www.symbol.com/product.php?productID=251</a>
Intermec CK61	<a href="http://www.intermec.com/products/cmptck61a/index.aspx">http://www.intermec.com/products/cmptck61a/index.aspx</a>
Intermec CK31	<a href="http://www.intermec.com/products/cmptck31/index.aspx">http://www.intermec.com/products/cmptck31/index.aspx</a>
Intermec CV60	<a href="http://www.intermec.com/products/cmptcv60/specs.aspx">http://www.intermec.com/products/cmptcv60/specs.aspx</a>
Psion Teklogix Workabout Pro	<a href="http://www.psionteklogix.com/assets/downloadable/WORKABOUT_PRO_1st-Gen_A4.pdf?ns=1">http://www.psionteklogix.com/assets/downloadable/WORKABOUT_PRO_1st-Gen_A4.pdf?ns=1</a>
Symbol (Motorola) MC3090	<a href="http://www.motorola.com/staticfiles/Business/Products/Mobile%20Computers/Handheld%20Computers/MC3000/_Documents/MC3000_SS_0308.pdf_5-14-08.pdf">http://www.motorola.com/staticfiles/Business/Products/Mobile%20Computers/Handheld%20Computers/MC3000/_Documents/MC3000_SS_0308.pdf_5-14-08.pdf</a>
Wireless Scale	Datasheet Link
Hobart Quantum wireless scale	<a href="http://www.hobartcorp.com/assets/specsheets/F-40159.pdf">http://www.hobartcorp.com/assets/specsheets/F-40159.pdf</a>

Table 13 lists the security modes test for the devices.

**Table 13** Single-Purpose HMD Supported Security Modes

Vendor	Device Type	Static WEP	WEP + .1x	WPA-PSK	WPA + .1x	WPA2-PSK	WPA2 +.1x
Symbol	MC3000	✓	✓ w/PEAP	✓	✓ w/PEAP	x	x
Symbol	MC50	✓	✓ w/PEAP	✓	✓ w/PEAP	x	x
Symbol	MC70	✓	✓ w/PEAP	✓	✓ w/PEAP	x	x
Symbol	MC9000	✓	✓ w/PEAP	✓	✓ w/PEAP	x	x
Symbol	PPT8800	✓	x	x	x	x	x
Symbol	PPT8100	✓	x	x	x	x	x
Symbol	VC5090	✓	✓ w/PEAP	✓	✓ w/PEAP	x	x
Symbol	MK2000	✓	x	x	x	x	x
Symbol	WT4090	✓	✓ w/PEAP	✓	✓ w/PEAP	x	x
Symbol	PDT6800	✓	x	x	x	x	x
Intermec	700 series	✓	✓ w/PEAP	✓	✓ w/PEAP	x	x
Intermec	CN2	✓	✓ w/PEAP	✓	✓ w/PEAP	✓	✓ w/PEAP
Intermec	CN3	✓	✓ w/PEAP	✓	✓ w/PEAP	✓	✓ w/PEAP
Intermec	CK31	✓	✓ w/PEAP	✓	✓ w/PEAP	✓	✓ w/PEAP
Intermec	CK60	✓	✓ w/PEAP	✓	✓ w/PEAP	x	x
Intermec	T2425	✓	x	x	x	x	x
Intermec	T2455	✓	✓ w/PEAP	✓	✓ w/PEAP	x	x
Intermec	CV60	✓	x	x	x	✓	✓ w/PEAP



**Table 13** Single-Purpose HMD Supported Security Modes (Continued)

Vendor	Device Type	Static WEP	WEP + .1x	WPA-PSK	WPA + .1x	WPA2-PSK	WPA2 +.1x
Teklogix	Workabout Pro	✓	✓ w/PEAP	✓	✓ w/PEAP	x	x
Teklogix	7530	✓	✓ w/PEAP	✓	✓ w/PEAP	✓	✓ w/PEAP
Teklogix	7535	✓	✓ w/PEAP	✓	✓ w/PEAP	✓	✓ w/PEAP
Vocollect	Talkman T5	✓	x	✓	x	x	x
Zebra	QL220	✓	x	✓	✓ w/PEAP	x	x
Zebra	RW220	✓	x	✓	x	x	x

Table 14 lists the mobility performance for the devices.

**Table 14** Single-Purpose HMD Supported Roaming Feature Enhancements

Vendor	Device Type	Fast Roaming	Standby Roaming	Load Balancing	PSP Support	Battery Boost
Symbol	MC3000	✓	✓	✓	✓	✓
Symbol	MC50	✓	✓	✓	✓	✓
Symbol	MC70	✓	✓	✓	✓	✓
Symbol	MC9000	✓	✓	✓	✓	✓
Symbol	PPT8800	✓	✓	✓	✓	✓
Symbol	PPT8100	✓	✓	✓	✓	✓
Symbol	VC5090	✓	✓	✓	✓	✓
Symbol	MK2000	✓	✓	✓	✓	✓
Symbol	WT4090	✓	✓	✓	✓	✓
Symbol	PDT6800	✓	✓	✓	✓	✓
Intermec	700 Series	✓	✓	✓	✓	✓
Intermec	CN2	✓	✓	✓	✓	✓
Intermec	CN3	✓	✓	✓	✓	✓
Intermec	CK31	✓	✓	✓	✓	✓
Intermec	CK60	✓	✓	✓	✓	✓
Intermec	T2425	✓	✓	✓	✓	✓
Intermec	T2455	✓	✓	✓	✓	✓
Intermec	CV60	✓	✓	✓	✓	✓
Teklogix	Workabout Pro	✓	✓	✓	✓	✓
Teklogix	7530	✓	✓	✓	✓	✓
Teklogix	7535	✓	✓	✓	✓	✓
Vocollect	Talkman T5	✓	✓	✓	✓	✓
Zebra	QL220	✓	✓	✓	✓	✓

## Commonly Deployed Voice HMDs

Table 15 lists some commonly deployed voice HMDs that have undergone interoperability testing with Aruba Networks and gives links to the datasheets for each device so you can gather more detailed information on each one.



**Table 15** Information Links for Certain Voice HMDs

Voice Handsets	Datasheet Link
Polycom SpectraLink 8020/8030 Wireless Telephones	<a href="http://www.polycom.com/usa/en/products/voice/wireless_solutions/wifi_communications/index.html">http://www.polycom.com/usa/en/products/voice/wireless_solutions/wifi_communications/index.html</a> <a href="http://www.polycom.com/common/documents/support/user/products/voice/SpectraLink_8020_8030_Getting_Started_Multi.pdf">http://www.polycom.com/common/documents/support/user/products/voice/SpectraLink_8020_8030_Getting_Started_Multi.pdf</a> <a href="http://www.polycom.com/common/documents/support/user/products/voice/SpectraLink_8020_8030_WT_User_Guide_SIP.pdf">http://www.polycom.com/common/documents/support/user/products/voice/SpectraLink_8020_8030_WT_User_Guide_SIP.pdf</a>
Alcatel/Lucent (ALU) NoE versions of SpectraLink phones with Wireless Multimedia (WMM)	<a href="http://www1.alcatel-lucent.com/com/en/appcontent/opgss/ENT_DATA_23030_610_Hndst_ds_EN_tcm228-1328681635.pdf">http://www1.alcatel-lucent.com/com/en/appcontent/opgss/ENT_DATA_23030_610_Hndst_ds_EN_tcm228-1328681635.pdf</a> <a href="http://www1.alcatel-lucent.com/com/en/appcontent/opgss/ENT_PHONES_IPTouch_610-310_WLAN_guide_0108_EN_tcm228-1383421635.pdf">http://www1.alcatel-lucent.com/com/en/appcontent/opgss/ENT_PHONES_IPTouch_610-310_WLAN_guide_0108_EN_tcm228-1383421635.pdf</a>
Vocera B2000 badges	<a href="http://www.vocera.com/downloads/Hardware_ds.pdf">http://www.vocera.com/downloads/Hardware_ds.pdf</a> <a href="http://www.vocera.com/downloads/UserGuide.pdf">http://www.vocera.com/downloads/UserGuide.pdf</a> <a href="http://www.vocera.com/downloads/ConfigGuide.pdf">http://www.vocera.com/downloads/ConfigGuide.pdf</a> <a href="http://www.vocera.com/downloads/InfrastructureGuide.pdf">http://www.vocera.com/downloads/InfrastructureGuide.pdf</a>
Cisco 7921g (a/b/g radios)	<a href="http://www.cisco.com/en/US/prod/collateral/voicesw/ps6788/phones/ps379/product_data_sheet0900aecd805e315d.pdf">http://www.cisco.com/en/US/prod/collateral/voicesw/ps6788/phones/ps379/product_data_sheet0900aecd805e315d.pdf</a>
Hitachi WIP5000A-E	<a href="http://www.wirelessip-phone.com/en/products/spec/index.html">http://www.wirelessip-phone.com/en/products/spec/index.html</a>
Nokia E51	<a href="http://europe.nokia.com/A4546209">http://europe.nokia.com/A4546209</a> <a href="http://nds1.nokia.com/files/support/apac/phones/guides/Nokia_E51_APAC_UG_en.pdf">http://nds1.nokia.com/files/support/apac/phones/guides/Nokia_E51_APAC_UG_en.pdf</a>
Nokia E61i	<a href="http://europe.nokia.com/A4344023">http://europe.nokia.com/A4344023</a> <a href="http://nds1.nokia.com/phones/files/guides/Nokia_E61i-1_UG_en.pdf">http://nds1.nokia.com/phones/files/guides/Nokia_E61i-1_UG_en.pdf</a>
Nokia E65	<a href="http://europe.nokia.com/A4344228">http://europe.nokia.com/A4344228</a> <a href="http://nds1.nokia.com/files/support/apac/phones/guides/Nokia_E65_APAC_UG_en.pdf">http://nds1.nokia.com/files/support/apac/phones/guides/Nokia_E65_APAC_UG_en.pdf</a>
Nokia E66	<a href="http://europe.nokia.com/A4344228">http://europe.nokia.com/A4344228</a> <a href="http://nds1.nokia.com/files/support/apac/phones/guides/Nokia_E66_APAC_UG_EN.pdf">http://nds1.nokia.com/files/support/apac/phones/guides/Nokia_E66_APAC_UG_EN.pdf</a>
Nokia E71	<a href="http://europe.nokia.com/A41146123">http://europe.nokia.com/A41146123</a> <a href="http://nds1.nokia.com/phones/files/guides/Nokia_E71-1_UG_en.pdf">http://nds1.nokia.com/phones/files/guides/Nokia_E71-1_UG_en.pdf</a>




Examples of  
Highly Mobile Devices

There are countless markets that require HMDs. An HMD is a device that is on the move for the majority of the day, coming to rest only at periodic intervals to be plugged into a battery recharging station. During periods of mobility, the device must stay connected to the wireless network as it roams from access point (AP) to access point. A few examples of HMDs are shown in [Table 16](#).


**Table 16** *HMD Devices Used in Various Markets*

Market	Example HMD Device	Description
Retail	 <p><b>SpectraLink SIP Phone</b></p>	Used for voice communications among employees and clerks in retail stores. SIP interoperability for Polycom SpectraLink 8000 Wi-Fi handsets allows the wireless telephones to work as extensions of the VoIP call server, and includes the most traditionally used business telephone features, including dial by extension, conferencing, call transfer, call forwarding, and caller ID. These call features are essential for mobile employees in the workplace where real-time telephone communication is mission-critical.
	 <p><b>Intermec CK60 Mobile Computer</b></p>	<p>Mobile workers both inside and outside the walls of a facility rely on highly mobile wireless mobile scanners to provide real-time visibility into business-critical information. Such scanners are designed to handle all day, every day use in nearly any environment to route salespeople, field repair personnel, and workers in the warehouse or on the retail floor to access, as well as capture and send, crucial business information in real time. Modern scanners contain the latest advances in mobile technology, have proven rugged construction, and support multi-mode wireless connectivity. All of this ensures accessibility to coworkers and business systems. Business processes are streamlined and errors are reduced, effectively increasing productivity—and profitability. Aruba's mobility solutions allow such scanners to roam both inside the facility and out, providing seamless wireless coverage.</p> <p>An example use case is a retail grocery store. Store clerks check store shelves several times a day to see if product supply is running low. If so, the clerk uses a wireless handheld device to scan the label on the shelf to initiate reordering. The clerk can also use a wireless device to make price changes on products by scanning the shelf label and adjusting the price, which then is reflected at the checkout register. New shelf price labels can be printed from a mobile printer.</p>

**Table 16** HMD Devices Used in Various Markets (Continued)

Market	Example HMD Device	Description
Warehouse	 <p><b>Smart Forklift</b></p>	<p>Many forklifts now have wireless PCs installed on them, making the location of items and empty storage space immediately visible. Forklifts have also been equipped with wireless printers that allow tickets to be printed that indicate where to pick up and deliver parts. Take the case of a large auto assembly plant where the company has invested heavily in wireless technology for their parts replenishment system. The system automatically updates every few minutes the exact whereabouts of every vehicle on-site, providing all the data needed to optimize work-in-process.</p> <p>So-called “Smart Forklifts” are increasingly used to transfer inventory to and from warehouse shelves. When an assembly-line operator takes the first component out of the last available line side stillage or container, the operator presses a wireless-enabled line side Smart Call button that triggers the initiation of a replenishment activity. This activity consists of printing a job ticket at a central “marketplace” in the parts storage facility at the plant. The forklift operator is notified and subsequently picks up the ticket and then picks up or delivers parts. The ticket data could also be sent directly to the wireless PC on the forklift itself, obviating the need to return to the marketplace to be allocated the next part replenishment job. In this manner, the company can make use of the forklift on the otherwise-empty return leg of the journey. For the whole operation to work properly, forklifts must remain connected to the wireless network as they traverse the assembly plant.</p>
	 <p><b>Vocera Voice Badge</b></p>	<p>Used by doctors, nurses, and staff to communicate with each other wirelessly within the hospital and through the PBX when outside of the hospital premises. The voice badge is a lightweight voice-activated device that uses an 802.11 (Wi-Fi) wireless network.</p> <p>The badge works in any location covered by the Medical Center's wireless network. The badge weighs about 2 ounces, is 4.2" long by 1.4" wide, and is worn around the neck with a lanyard, clipped to a jacket lapel or shirt pocket, or attached to a belt, with an ear piece and microphone attached to it.</p> <p>The badges are voice activated and recognize speech patterns. The communication system can easily learn heavily accented speech patterns. Users are set up in functional groups so that organizations can be contacted instead of having to know the names of each individual.</p>
Healthcare	 <p><b>Tablet PC</b></p>	<p>The healthcare industry makes use of tablet PCs. Some of the newer proposed highly mobile tablet PC devices include a 10.4-inch, 1,024 × 768 touchscreen as well as Wi-Fi and a digital camera. The goal is to use these HMDs to allow healthcare personnel to streamline data entry and better monitor wounds and healing. These types of devices are targeted to read both barcodes and RFID tags to prevent medication errors and confirm staff and patient ID. In addition, doctors use the tablets to share results with patients. They can display X-rays and cardiac catheter images, for example. The imagery helps doctors when they discuss the course of treatment with patients. Tablet PCs are very mobile and move with the doctors as they make their rounds. As a result, they often roam among APs and must stay connected to the network.</p>

**Table 16** *HMD Devices Used in Various Markets (Continued)*

Market	Example HMD Device	Description
Healthcare (continued)	 <p><b>Workstation on Wheels</b></p>	<p>Workstations on Wheels (WOWs) are used to hold wireless patient monitoring devices as well as laptop and other wireless devices needed to accompany patients as they are transported from place to place.</p> <p>With the proliferation of electronic medical records (EMRs) and maturing wireless technologies, WOWs have become the predominant point-of-care (POC) mobile device form factor. According to HIMSS Analytics (2006), more than 75% of U.S. hospitals report some level of WOW adoption.</p> <p>As EMRs proliferate, increasing numbers of physicians and allied health professionals, as well as patients and their families, are leveraging evolving WOW technology.</p> <p>Next generation WOWs are being enhanced to improve POC access, functionality, technology integration, and telehealth. New capabilities include integration with software that captures data from diverse patient monitors and medical devices (IV smart pumps and ventilators) for seamless transfer of information into the patient's EMR, with appropriate validation by clinicians, where it immediately is available for physician and care team review.</p>



This appendix presents key platform capacity ceilings, transaction rates, and roaming time values. It is intended to assist the wireless architect in validating the scalability of a design prior to deployment.

## Network Design Scaling Considerations

The following platform specifications indicate the maximum supported values for various key design parameters .

**Table 17** *Mobility Controller Product Line Matrix*

Features	MC-200	MC-800		MC-2400	MMC-3000 Series		
		MC-804	MC-800		MMC-3200	MMC-3400	MMC-3600
Max number of campus-connected APs per controller	6	4	16	48	32	64	128
Max number of Remote APs per controller	6	4	16	48	128	256	512
Max number of users per controller	100	256	255	768	512	1,024	2,048
MAC Addresses	4,096	4,096	4,096	4,096	64,000	64,00	64,000

**Table 18** *Controller Specifications*

Features	MMC-6000	MMC-6000 Supervisor Card	
		SC-256-C2	M3
Max number of campus-connected APs per controller	2,048	256	512
Max number of Remote APs per controller	8,192	256	2,048
Max number of users per controller	32,768	4,096	8,192
MAC Addresses	256,000	64,000	64,000

[Table 19](#) indicates the maximum number of APs and users that can be managed by a single master controller. For layer 3 (IP Mobility) deployments, the mobility domain boundaries should be sized with these values in mind.

**Table 19** *Maximum Number of APs and Users per Master Controller Model*

Master	Maximum APs	Maximum Users
M3 Blade/MMC-3600	4,500	15,000
Supervisor II Blade	3,000	10,000
MMC-3400	2,250	7,500
MMC-3200	1,500	4,500

## Controller Scaling Considerations

Table 20 provides validated performance limits for Aruba's two chassis-based supervisor blades.

**Table 20** Validated Performance Limits for Aruba Chassis-based Supervisor Blades

Metric	SupervisorII	M3
Maximum Authentications per Second (Open SSID)	360	720
Maximum Number of User Roles	300	500
Maximum Concurrent Active Bandwidth Contracts	200	512
Maximum Number of DHCP leases per second	18	26
Maximum Number of DHCP leases	1200	1400
Maximum Number of Multicast Groups	1,024	1,024

## Data Device Roaming Considerations

Aruba has measured L2 and L3 roaming times for HMD data clients in the ranges shown below. Exact performance will vary based on device processor speed, radio type, network utilization and other factors.

**Table 21** Data HMD Clients

	With OKC	No OKC	
	Intra-Controller	Intra-Controller	Inter-Controller
Layer 2	19 ms	76 ms	83 ms
Layer 3	17 ms	185 ms	229 ms

## Voice Device Roaming Considerations

Aruba has conducted roaming tests with commonly deployed voice devices. The following values are for informational purposes only to assist the wireless architect in making design decisions. Values are approximate, and your results may vary based on the client device processor speed, radio type, network utilization and other factors.

Wireless roaming with 802.1X involves several parts where delay can be attributed to device or infrastructure activity:

- The wireless device has an internal algorithm that will cause it to start looking for a better AP once its driver conditions are met.
- The device needs to find an AP by sending 802.11 Probe Requests with its configured SSID to each supported channel for that phy-type it is scanning. Once the device finds an AP, it will build a list of BSSIDs, Signal Strength, and other possible information to use to connect to the "best" AP via an 802.11 Auth frame with its gathered information.
- If the AP sees the 802.11 Auth frame it will respond with the 802.11 Auth frame with status "success."
- The 802.11 exchange will continue with an 802.11 Association Request from the client and an 802.11 Association Response from the AP if the client 802.11 requested information meets the AP configuration settings.



- An 802.1X authentication exchange between the device and the AP then begins.
- After there is a successful 802.1x authentication exchange, the encryption keys are completely primed on the device and network access is allowed.

Aruba strongly recommends the use of PMK caching. In the following three tables, we measured the delay between the client sending the first 802.11 Auth packet until the last key was exchanged. 802.1X roaming times when PMK caching is not used depend directly on handset device processor speed and driver algorithms.

**Table 22** Cisco 7921g Roaming Performance (WPA2-AES Enterprise, EAP-PEAP)

	PMK Cached			
	Intra-Controller Intra-VLAN	Intra-Controller Inter-VLAN	Inter-Controller Intra-VLAN	Inter-Controller Inter-VLAN
<b>Min.</b>	56 ms	62 ms	63 ms	56 ms
<b>Max.</b>	86 ms	67 ms	70 ms	62 ms
<b>Avg.</b>	68 ms	64 ms	66 ms	59 ms

**Table 23** Nokia E71 Roaming Performance (WPA2-AES Enterprise, EAP-PEAP)

	PMK Cached			
	Intra-Controller Intra-VLAN	Intra-Controller Inter-VLAN	Inter-Controller Intra-VLAN	Inter-Controller Inter-VLAN
<b>Min.</b>	19 ms	17 ms	17 ms	18 ms
<b>Max.</b>	24 ms	22 ms	24 ms	23 ms
<b>Avg.</b>	21 ms	19 ms	19 ms	20 ms

**Table 24** Vocera Roaming Performance (WPA-PSK)

	Intra-Controller Intra-VLAN	Intra-Controller Inter-VLAN	Inter-Controller Intra-VLAN	Inter-Controller Inter-VLAN
<b>Min.</b>	50 ms	50 ms	44 ms	54 ms
<b>Max.</b>	58 ms	55 ms	53 ms	58 ms
<b>Avg.</b>	55 ms	52 ms	50 ms	56 ms



## Contacting Aruba Networks

Web Site Support	
Main Site	<a href="http://www.arubanetworks.com">http://www.arubanetworks.com</a>
Support Site	<a href="https://support.arubanetworks.com">https://support.arubanetworks.com</a>
Software Licensing Site	<a href="https://licensing.arubanetworks.com/login.php">https://licensing.arubanetworks.com/login.php</a>
Wireless Security Incident Response Team (WSIRT)	<a href="http://www.arubanetworks.com/support/wsirt.php">http://www.arubanetworks.com/support/wsirt.php</a>
Support Emails	
• Americas and APAC	<a href="mailto:support@arubanetworks.com">support@arubanetworks.com</a>
• EMEA	<a href="mailto:emea_support@arubanetworks.com">emea_support@arubanetworks.com</a>
WSIRT Email Please email details of any security problem found in an Aruba product.	<a href="mailto:wsirt@arubanetworks.com">wsirt@arubanetworks.com</a>

Telephone Support	
Aruba Corporate	+1 (408) 227-4500
FAX	+1 (408) 227-4550
Support	
• United States	+1-800-WI-FI-LAN (800-943-4526)
• Universal Free Phone Service Numbers (UIFN):	
■ Australia	Reach: 11 800 494 34526
■ United States	1 800 9434526 1 650 3856589
■ Canada	1 800 9434526 1 650 3856589
■ United Kingdom	BT: 0 825 494 34526 MCL: 0 825 494 34526
■ Japan	IDC: 10 810 494 34526 * Select fixed phones IDC: 0061 010 812 494 34526 * Any fixed, mobile & payphone KDD: 10 813 494 34526 * Select fixed phones JT: 10 815 494 34526 * Select fixed phones JT: 0041 010 816 494 34526 * Any fixed, mobile & payphone
■ Korea	DACOM: 2 819 494 34526 KT: 1 820 494 34526 ONSE: 8 821 494 34526
■ Singapore	Singapore Telecom: 1 822 494 34526
■ Taiwan (U)	CHT-I: 0 824 494 34526

## Telephone Support

■ Belgium	Belgacom: 0 827 494 34526
■ Israel	Bezeq: 14 807 494 34526 Barack ITC: 13 808 494 34526
■ Ireland	EIRCOM: 0 806 494 34526
■ Hong Kong	HKTI: 1 805 494 34526
■ Germany	Deutsche Telekom: 0 804 494 34526
■ France	France Telecom: 0 803 494 34526
■ China (P)	China Telecom South: 0 801 494 34526 China Netcom Group: 0 802 494 34526
■ Saudi Arabia	800 8445708
■ UAE	800 04416077
■ Egypt	2510-0200 8885177267 * within Cairo 02-2510-0200 8885177267 * outside Cairo
■ India	91 044 66768150