
VALIDATED REFERENCE DESIGN



VERY HIGH-DENSITY 802.11ac NETWORKS

Engineering and
Configuration Guide

Version 1.0

Chuck Lukaszewski, CWNE #112



Copyright

© 2015 Aruba Networks, Inc. All rights reserved. Aruba Networks®, Aruba Networks™ (stylized), People Move Networks Must Follow®, Mobile Edge Architecture®, RFProtect®, Green Island®, ClientMatch®, Aruba Central®, Aruba Mobility Management System™, ETips™, Virtual Intranet Access™, Aruba Instant™, ArubaOS™, xSec™, ServiceEdge™, Aruba ClearPass Access Management System™, AirMesh™, AirWave™, Aruba@Work™, Cloud WiFi™, Aruba Cloud™, Adaptive Radio Management™, Mobility-Defined Networks™, Meridian™ and ArubaCare™ are trademarks of Aruba Networks, Inc. registered in the United States and foreign countries. Aruba Networks, Inc. reserves the right to change, modify, transfer or otherwise revise this publication and the product specifications without notice. While Aruba Networks uses commercially reasonable efforts to ensure the accuracy of the specifications contained in this document, Aruba Networks will assume no responsibility for any errors or inaccuracies that may appear in this document.

Open Source Code

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

http://www.arubanetworks.com/open_source

Legal Notice

ARUBA DISCLAIMS ANY AND ALL OTHER REPRESENTATIONS AND WARRANTIES, WEATHER EXPRESS, IMPLIED, OR STATUTORY, INCLUDING WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, NONINFRINGEMENT, ACCURACY AND QUET ENJOYMENT. IN NO EVENT SHALL THE AGGREGATE LIABILITY OF ARUBA EXCEED THE AMOUNTS ACUTALLY PAID TO ARUBA UNDER ANY APPLICABLE WRITTEN AGREEMENT OR FOR ARUBA PRODUCTS OR SERVICES PURSHASD DIRECTLY FROM ARUBA, WHICHEVER IS LESS.

Warning and Disclaimer

This guide is designed to provide information about wireless networking, which includes Aruba Network products. Though Aruba uses commercially reasonable efforts to ensure the accuracy of the specifications contained in this document, this guide and the information in it is provided on an "as is" basis. Aruba assumes no liability or responsibility for any errors or omissions.

ARUBA DISCLAIMS ANY AND ALL OTHER REPRESENTATIONS AND WARRANTIES, WHETHER EXPRESSED, IMPLIED, OR STATUTORY, INCLUDING WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, NONINFRINGEMENT, ACCURACY, AND QUIET ENJOYMENT. IN NO EVENT SHALL THE AGGREGATE LIABILITY OF ARUBA EXCEED THE AMOUNTS ACTUALLY PAID TO ARUBA UNDER ANY APPLICABLE WRITTEN AGREEMENT OR FOR ARUBA PRODUCTS OR SERVICES PURCHASED DIRECTLY FROM ARUBA, WHICHEVER IS LESS.

Aruba Networks reserves the right to change, modify, transfer, or otherwise revise this publication and the product specifications without notice.



1344 CROSSMAN AVENUE | SUNNYVALE, CALIFORNIA 94089
1.866.55.ARUBA | T: 1.408.227.4500 | arubanetworks.com

Table of Contents

Chapter EC-1: Introduction	8
Chapter EC-2: Estimating System Throughput	9
System Throughput vs. Channel Throughput vs. Device Throughput	9
Total System Throughput	10
Channel Throughput	11
Device Throughput	12
Single-client Throughput	12
Setting Proper Expectations	13
Converting Data Rates to Throughput	14
Data Rate is Assumed to be Maximized	15
Multi-client Throughput	15
Multi-client Throughput Test Results	15
Using Lab Results for VHD Capacity Planning	17
Modeling the Spatial Stream Mix of VHD Areas	17
Estimation Process for System Throughput	19
Step 1 – Select Channel Count	19
Step 2 – Estimate Unimpaired Multi-client Throughput	26
Step 3 – Choose and Apply Impairment Factor	26
Step 4 – Select Spatial Reuse Factor	28
Step 5 – Calculate Total System Throughput	30
Estimation Process for Per-Device Throughput	30
Examples	31
Example #1 – Small Auditorium	31
Example #2 – Indoor Arena in China	32
Example #3 – Outdoor Stadium in USA	33
Comparing Examples to Aruba VHD Lab Results	35
Conclusions	36
Chapter EC-3: Airtime Management	37
Maximizing Capacity Through Good Configuration	37
Good Client Distribution	38
Send Transmissions in Parallel by Using Only 20-MHz Channels	38
Use All 5-GHz Channels (Except 144)	41
Ensure a Good Channel Plan	42
Use Power Differentials to Self-Steer Clients to 5-GHz	42
Use ClientMatch and 802.11v BSS Transition for Steering and Load Balancing	43
Enable 802.11k Radio Measurements	44
Use QBSS Load to Self-Steer Clients to Less Busy APs	44

Maximize Data Rates	45
Enable VHT Rates on 2.4-GHz	45
Maximize Rate of 802.11a and 802.11n Data Frames by “Trimming” Low Rates	45
Increase Data Rate of Control and Management Frames	46
Disable 802.11b Rates and Protection Mode	47
Use a High Beacon Rate	47
Enable Dynamic Multicast Optimization (DMO)	49
Enable Multicast-Rate-Optimization	50
Eliminate Unnecessary Transmissions	51
Eliminate SSIDs and Beacons	51
Drop Broadcast and Multicast to Block “Chatty” Protocols	51
Do Not Use IPv6	52
Use AirGroup If mDNS or DLNA Services Are Required	52
Enable IGMP Proxy	52
Deny Inter-User Bridging	53
Set a Local Probe Response Threshold	53
Increase A-MSDU to 3 and Enable End-to-End Jumbo Frames	53
Maximize Channel Accessibility	54
Use Quality of Service	54
Use Airtime Fairness	54
Do Not Use Bandwidth Contracts	55
Use the Cell Size Reduction Feature to Filter Low-SINR CCI	56
Chapter EC-4: Channel and Power Plans	59
How VHD Channel Requirements Differ from Conventional WLANs	59
Facilities with Multiple Adjacent VHD Zones	59
Stadiums and Arenas	60
Responding to Interference	60
Types of Channel Plans	61
Choosing Dynamic vs. Static Channel Plans	61
Requirements of Good VHD Channel Plan	62
Should You Disable “Surplus” 2.4-GHz Channels?	63
How VHD Power Requirements Differ from Conventional WLANs	63
Facilities with Multiple Adjacent VHD Zones	64
Stadiums and Arenas	64
Requirements of Good VHD Power Plan	64
Power Plans Should Deliver Client SINR Equal or Greater than 30 dB	64
Use the Absolute Minimum EIRP Possible in Each VHD Area	65
Power Should Be Measured by 1SS Devices	65
Use a 6 – 9 dB Lower Power on the 2.4-GHz Band	65
Use a Static EIRP Setting Based on Post-Install Tuning	65

How Aruba Adaptive Radio Management Makes Decisions	66
Background Scanning	66
Radio Metrics	66
The Regulatory Domain Profile	67
ARM Profile	68
Configuring Dynamic and Static Channel Plans	68
Configuring the ARM Assignment Mode	68
Initial Channel Assignment for Static Plans	68
Configuring Global and Local Channel Plans	69
Setting Up a Regulatory Domain Profile	69
Disabling 40-MHz and 80-MHz Channel Widths	71
Setting Up a Static 1, 6, 11 Channel Plan for a Stadium	72
Setting Up Other Local Channel Plans	74
Configuring Repeating and Non-Repeating Channel Plans	75
Configuring the Power Plan	75
Chapter EC-5: SSIDs, Authentication, and Security	76
Overview	76
SSID Strategy for VHD Networks	76
SSID Summary	77
Open Guest Network with Captive Portal	78
Converge Secure Users on to One SSID	78
Role and Policy Strategy for VHD Networks	78
Guest SSID Design	81
Captive Portal	81
Certificates	82
MAC Caching	82
SSID and Virtual AP Configuration	82
Secure SSID Design	82
Compatibility Requirements	83
Preshared Key Security Risks	83
Replicating the PSK Experience with 802.1X	84
Zero Tolerance for Third-Party APs	84
SSID and Virtual AP Configuration on Controller	85
Dynamic Authorization Profiles on ClearPass Policy Manager	85
MAC Address Authentication	87
ClearPass Policy Design	87
Offload SSID Design	88
Passpoint	88
Personal Hotspots	88

Post-Deployment Site Surveys	89
VHD Network Hardening	89
Links to ClearPass Technical Documents	90
Chapter EC-6: Video Streaming	91
Dimensioning Video Usage	91
Dimensioning Process	91
Computing a Bitrate Table for a Stadium	92
Computing a Bitrate Table for a Lecture Hall	93
Multicast vs. Unicast	94
Video Quality vs. System Capacity	96
Configuring Unicast Video Streaming	96
Ensure End-to-End QoS Marking	97
Configure Remarking ACLs	97
Verify MTU Size	97
Ensure Video-Aware ARM is Enabled	97
Configuring Multicast Video Streaming	97
General Multicast Configuration	98
Set a Fixed Multicast Rate for Video Traffic	98
Use Forward Error Correction	98
Video Upload	99
Further Assumptions	99
Chapter EC-7: Configuration Summary	100
Virtual AP Profile	100
SSID Profile	101
HT-SSID Profile	102
ARM Profile	102
Dot11a Radio Profile	103
Dot11g Radio Profile	103
Regulatory Domain Profile	104
Traffic Management Profile	104
Dot11k Profile	104
Radio Resource Measurement IE Profile	104
AAA Authentication Server Group	105
VLAN Interface	105
Hardening Checklist	105

Appendix EC-A: Worldwide 5-GHz Channel Availability as of March 1, 2015	106
Appendix EC-B: 802.11ac Data Rate Table	107
Appendix EC-C: DFS Surveys and Operating Rules	108
How to Conduct a DFS Survey	108
Behavior of 5-GHz Client Devices in Presence of Radar	109
DFS Fact vs. Fiction	110

Chapter EC-1: Introduction

Welcome to the Engineering and Configuration guide of the Aruba Very High Density (VHD) Validated Reference Design (VRD). The previous guide (Planning) explains what a VHD network is, presents a structured methodology for dimensioning an end-to-end system, explains how to choose APs and antennas, and introduces the three basic radio coverage strategies that can be used. That guide is aimed at a wide audience.




Figure EC1-1 Organization of the Very High Density VRD

This guide is far more technical. It is aimed at the engineers of our customers and partners. After reading this guide, you should be able to:

- Predict usable bandwidth in a VHD environment
- Configure various Aruba features that optimize airtime availability
- Design and implement dynamic and static channel plans with ARM
- Architect an airtime-efficient SSID design
- Explain to others why the SSID design must follow this guidance
- Properly configure the controller for unicast or multicast video
- Understand how to dimension video consumption and its impact on other services

This Engineering and Configuration guide is aimed at engineers who simply need to know the recipe that they should follow to get a VHD network up and running. This guide does not attempt to explain the theory of operations behind the recommendations provided.

For WLAN architects, we recommend that you continue with the Theory guide after you have fully grasped the content of this guide.

All readers should also review the appropriate Scenario document for their particular high-density use case.

Chapter EC-2: Estimating System Throughput

After the work done in the Planning guide, the very high density (VHD) network has been dimensioned and a bill of materials has been created. A high-level radio design is being developed.

At this stage, the most common question customers ask about VHD wireless LANs is *“how much guaranteed throughput will the network provide?”*

This question is very difficult to answer unless you actually build the system. The shared nature of the wireless medium and the potentially very large number of devices attempting to share it make accurate estimates difficult. Critical factors such as the level of RF spatial reuse and RF interference can be studied by spectrum analyzers or protocol analyzers. But we have no standard way to apply the resulting data to capacity estimates.

Nevertheless, growing numbers of customers who purchase a VHD system define their requirements in terms of minimum per-seat or per-device throughput. Some customers even attempt to contractually guarantee such minimums. Often the driver for this is a software application like video streaming that needs a certain bitrate for a good experience. So a method must be found to provide a good-faith estimate, based on the unique particulars of each facility.

The question of guaranteed throughput clearly demarcates the largely numbers-driven, dimensioning and selling process from the mostly technical engineering and configuration process. Therefore, we begin the Engineering and Configuration guide by addressing guaranteed throughput first. In this chapter, you learn how to estimate the total aggregate capacity that may be available to users of the system. Then we cover how to convert this total system throughput value to a per-device estimate.

System Throughput vs. Channel Throughput vs. Device Throughput

Aruba has installed numerous high-density deployments of all sizes. This experience shows that the best approach to this question is to start by estimating the gross system capacity, and then work backward to the device level.

The total capacity of a VHD service area essentially is fixed, notwithstanding external interference. Aruba calls this the “total system throughput” of a VHD network. System throughput is the sum of the capacities of each individual channel, as shown in [Figure EC2-1](#). Every channel is assumed to have the same potential throughput unless a known interference source exists in that particular frequency range.

One common misconception about high density WLANs is that one can increase capacity by adding APs. This is false. The number of access points (APs) has nothing to do with the capacity of the system, unless you have RF measurements that prove that RF spatial reuse is possible in your facility. APs normally have to fight for airtime just like any user device. This is shown graphically in [Figure EC2-1](#) by the presence of multiple APs in each channel block.




Figure EC2-1 Defining Throughput Terminology for a VHD Wireless LAN

Within each channel, Wi-Fi® clients and APs divide the available airtime amongst themselves in real time. When a client gains control of the channel, the throughput that it can achieve depends on its particular capabilities and the quality of its radio link to the AP. The 802.11ac MAC protocol limits the maximum time a single device can reserve the channel to 5.4 ms. However, the vast majority of transmissions are simple control frames or TCP acknowledgments that are far shorter than this.

Total System Throughput

Total system throughput (TST) can be thought of as the total load that is generated by the VHD network through the WAN interface(s) of the Internet edge router(s). TST is estimated with this equation:

$$TST = Channels * Average Channel Throughput * Reuse Factor \quad (1)$$

Where:

- Channels = Number of channels in use by the VHD network
- Average Channel Throughput = Weighted average goodput that is achievable in one channel by the expected mix of devices for that specific facility
- Reuse Factor = Number of RF spatial reuses possible. For all but the most exotic VHD networks, this is equal to 1, which means no spatial reuse.

These variables are discussed in depth shortly. For the moment, consider this simple example to get a feel for the formula. Assume a 5-GHz deployment with 9 channels, an average channel throughput of 50 Mbps, and no RF spatial reuse. The TST would be 450 Mbps as follows:

$$TST = 9 \text{ Channels} * 50 \text{ Mbps} * 1 \text{ Reuse} = 450 \text{ Mbps}$$

The TST is the highest load you would ever expect to see on the WAN router uplink. Therefore, the WAN link must be dimensioned larger than this. Later in this chapter, we will work through several detailed examples.



Remember that the controller uplink is slightly more than $TST * 2$ due the hairpin turn that all WLAN traffic must take through the controller after it exits the AP tunnel. A material amount of tunnel encapsulation overhead is included also.

Why Not Count APs Instead of Channels?

Notice that formula (1) mentions channels but not APs. Adding APs almost never increases capacity in VHD areas, due to co-channel interference. This concept is widely misunderstood. So why use more than one AP per channel? For two main reasons:

- **To meet the associated device capacity (ADC) target:** Each AP radio serves just 150 users (or whatever value you chose in the dimensioning step), so a VHD system often requires many more APs than channels to ensure that all users can be associated to the network.
- **To ensure a high SINR throughout the facility:** Even if one central AP could handle the load, distant clients would connect at slower rates than closer clients due to lower signal-to-interference-and-noise ratio (SINR). When this occurs, it slows down the whole channel. Using more APs closer to clients keeps data rates high and helps transmissions finish faster, but the cost is higher CCI.



Most WLAN engineers and vendors use the term signal-to-noise ratio (SNR). In VHD areas with high levels of ongoing interference, chances are good that interference is higher than the noise floor. So SINR is a more accurate term.

Channel Throughput

Every Wi-Fi channel has a specific limit to the capacity it can offer. Broadly speaking, this limit is determined by the width of the channel and the number of spatial streams that are supported by the clients that attempt to use it. Channel width is 20-MHz, 40-MHz, 80-MHz, or 160-MHz. The 802.11ac standard provides for up to eight spatial streams, though most client devices only support one or two.

In addition, channel throughput depends heavily on the number of stations that attempt to use the channel simultaneously. **Capacity is actually lost due to collisions and MAC-layer inefficiencies as more devices contend for access.** Therefore, it is important to distinguish between single-client throughput and multiclient throughput.

Single-client throughput is basically the best number you get during a speed test on a clean channel with no other users present. The only real use of this value is to estimate how fast the WLAN should be when the high-density area is empty. This number is important to the test methodology that is used during system acceptance.

Multi-client throughput is the weighted average goodput that is achievable in one channel by the expected mix of devices in a particular VHD area. This throughput is generally much **less** than the single-client throughput. The multi-client throughput number gives you an accurate idea of how the network will perform when the VHD area is full of users attempting to access the WLAN.

Channel throughput can be *further reduced* by many factors including misbehaving client devices, CCI, ACI and non-Wi-Fi interference. In general, these factors are referred to as “impairments.” The term “average channel throughput” in formula (1) is meant to capture all of these effects for a given environment. For example, if you know from a spectral analysis that some channels have significant ongoing interference you must discount your multi-client throughput accordingly.

You will learn much more about estimating single-client throughput, multi-client throughput, and impairments in the coming pages.

Device Throughput

Device throughput is the average produced by dividing the TST by the number of devices expected to use the VHD system at the exact same moment in time. This equation is written as follows:

$$\text{Average Device Throughput} = \frac{\text{Total System Throughput}}{\text{Instantaneous User Count}} \quad (2)$$

The instantaneous user count is the number of devices that actually attempt to use the channel. This number is always significantly lower than the number of associated 802.11 devices. The instantaneous user count depends primarily on the duty cycle of the applications running on the devices.

Device throughput varies from channel to channel and from moment to moment based on many factors. The most obvious factor is fluctuation in the number of devices that want to transmit. If most associated devices have nothing to send, then device throughput rises dramatically. If many devices are attempting to send, device throughput can fall just as quickly. In either case, the devices are dividing the channel capacity among themselves in real time.

Most customers who purchase a VHD network state their expectations in terms of per-user or per-device throughput. It is common to hear requirements for a guaranteed per-device SLA such as 512 Kbps or 1 Mbps. This requirement is especially true if video streaming is to be provided, for which a specific video bitrate must be selected.



Generally, it is not possible to guarantee a specific per-device value in a VHD system. The best that can be achieved is to show that for a given VHD configuration, with a given number of channels, with a given client mix, with a given duty cycle, with a clean channel, that the average device throughput will be approximately some number of Kbps or Mbps. In VHD areas with ADC values over 1,000 it is common to talk in terms of Kbps.

Single-client Throughput

Later in the chapter, the TST formula is analyzed in detail, and you will work through several examples. First, we must introduce some key concepts and terminology that provide the building blocks needed to understand VHD capacity planning.

The first building block is to understand the potential performance of a single Wi-Fi client. When you know how to estimate single-client throughput, we can proceed to the more complex multi-client case.

$$\text{Single Client Throughput ("Goodput")} = \text{Peak Data Rate} * (1 - \text{Protocol Overhead\%}) \quad (3)$$

The peak data rate is the fastest Layer 1 (PHY) modulation and coding scheme (MCS) data rate that a particular client device is capable of achieving. The peak data rate depends on the width of the channel (20-MHz, 40-MHz, or 80-MHz) and the number of spatial streams that client supports. A complete table of 802.11ac data rates is found in [Appendix EC-B: 802.11ac Data Rate Table](#).

Protocol overhead includes Layer 2 (MAC) and Layers 3 and 4. With 802.11ac, we recommend you use a constant of 25% for MAC+TCP, or 20% for MAC+UDP.

Setting Proper Expectations

When you know the single-client throughput, you help ensure that all parties have the same understanding about the expected level of performance during speed tests.

Sometimes when engineers do VHD capacity planning, they fail to consider the channel width or the radio capabilities of the expected client device population. Currently shipping 802.11ac APs with three spatial stream (3SS) support are capable of up to 1.3 Gbps data rates in an 80-MHz channel. So, it is not unusual to hear customers using such large figures in discussing their VHD expectations.

But such high rates are absolutely incorrect for VHD planning for several reasons.

1. VHD areas should never be designed for peak single-client burst rate. VHD areas are designed to provide a low, common throughput like 512 Kbps or 1 Mbps to all clients. While it may occasionally be possible to hit the peak if the network is not busy, the baseline assumption for any high-density network is that the channel is very congested and average device throughput is much lower than the peak rate.
2. VHD areas should use only 20-MHz channel widths instead of 80-MHz widths. This narrower channel width brings down the peak data rate dramatically (from 1.3 Gbps to just 86.7 Mbps for a 1SS 802.11ac smartphone).
3. Most client devices in VHD areas are only 1SS or 2SS ([Table EC2-1](#)). 1SS or 2SS devices are expected to be norm in the future due to device size and battery power constraints.

Table EC2-1 Spatial Stream Capabilities of Common Devices at Publication

Type	Device Make and Model	Radio	1SS	2SS	3SS
Smartphone	iPhone 5 (All models)	11n	X		
	iPhone 6 (All models)	11ac	X		
	Samsung Galaxy S4	11ac	X		
	HTC One	11ac	X		
	LG G3	11ac	X		
	Windows Phone	11n	X		
	Samsung Galaxy S5	11ac		X	
Tablet	Apple iPad Air	11n		X	
	Apple iPad Air 2	11ac		X	
	Samsung Galaxy Tab 7.0 & 10.1	11n	X		
	Samsung Galaxy Tab Pro 12.2	11ac		X	
	Microsoft Surface Pro 3	11ac		X	

Table EC2-1 Spatial Stream Capabilities of Common Devices at Publication (Continued)

Type	Device Make and Model	Radio	1SS	2SS	3SS
Laptop / Netbook	Chromebook	11n		X	
	Apple MacBook Air	11ac		X	
	Apple MacBook Pro	11ac			X
	Intel 4965agn	11n		X	
	Intel 5300agn	11n			X
	Intel 6300agn	11n			X
	Intel 7260	11ac		X	

In addition, legacy devices such as 802.11n devices do not support the latest 256-QAM data rates introduced in 802.11ac. These devices will continue to be a significant percentage of the device population for many years. [Table EC2-2](#) shows the maximum MCS rate possible for both legacy 802.11n and newer 802.11ac clients in a 20-MHz channel width.

Table EC2-2 Maximum PHY Data Rate in a 20 MHz Channel

Device Radio Type	Peak MCS	1SS	2SS	3SS	4SS
802.11n (HT)	MCS 7	72.2 Mbps	144.4 Mbps	216.7 Mbps	n/a
802.11ac (VHT)	MCS 8** / 9	86.7 Mbps	173.3 Mbps	288.9 Mbps	346.7 Mbps



MCS 8 is the highest data rate in a VHT20 channel for 1SS, 2SS, and 4SS clients. 3SS clients support MCS 9.

So it is critically important to set proper expectations when thinking about VHD networks. Proper expectations help to avoid misunderstandings if someone runs a speed test and sees lower numbers than expected.

Converting Data Rates to Throughput

We can apply formula (3) and the protocol constants to each of the cells of [Table EC2-2](#) to convert from PHY data rate to single-client throughput.

Table EC2-3 Usable TCP Goodput in a 20 MHz Channel

Device Radio Type	Peak MCS	1SS	2SS	3SS	4SS
802.11n (HT)	MCS 7	54 Mbps	108 Mbps	162 Mbps	n/a
802.11ac (VHT)	MCS 8** / 9	65 Mbps	130 Mbps	217 Mbps	260 Mbps

These values are the absolute best case numbers you would expect to see in a speed test, when using an AP of equal or greater capability. No one else can be using the channel during the test.

Data Rate is Assumed to be Maximized

This chapter assumes that the RF coverage strategy ensures that the highest MCS supported by clients generally can be achieved. The beginning of [Chapter P-3: RF Design](#) of the *Very High-Density 802.11ac Networks Planning Guide* states that the first responsibility of the VHD architect is:

"To ensure that a minimum 25dB – 30dB SINR is available everywhere in the service area (after accounting for signal loss due to human bodies)"

If this requirement is met, then we safely can assume that the peak MCS value is available to each client.

If the minimum SINR is lower than 25dB for some reason, then the performance values discussed in this chapter for single-client and multi-client throughput must be reduced proportionately. Talk to your local Aruba systems engineer for assistance in this case.

Multi-client Throughput

Though single-client throughput can be computed from data rates, multiclient throughput must be tested with real clients in open air. Real tests are required because the capacity of the channel actually decreases as the number of clients increases. The total combined throughput of 25 devices that compete for airtime will be less than the single-client value. In turn, 50 clients achieve even less aggregate throughput than 25, and so on.

The main reason for this decrease in capacity is MAC-layer overhead that consumes a larger and larger share of the airtime as more stations contend for the medium. For the purposes of this guide, that fact is all you need to know to produce a system throughput estimate. WLAN architects should read [Chapter T-3: Understanding Airtime](#) in the *Very High-Density 802.11ac Networks Theory Guide*, which provides details about the underlying mechanics of this phenomenon.

Multi-client Throughput Test Results

[Figure EC2-2](#) shows a client scaling test result from the Aruba VHD lab. Client scaling tests measure performance with increasing numbers of real clients in open air. In this case, we started with 10 clients, and then tested 25, 50, 75, and finally 100 clients. The traffic type was TCP and we are showing bidirectional results in which each station is both sending and receiving data at the same time (which are the most difficult multi-client tests). Data was taken in a clean VHT20 channel.




Figure EC2-2 TCP Test Results for 100 Simultaneous 802.11ac Devices

The maximum throughput changes with spatial stream capability, so we ran the entire test separately for 1SS, 2SS, and 3SS devices. The Aruba VHD testbed has 100 of each type of devices as listed in [Table EC2-4](#).

Table EC2-4 Device Makeup of Aruba VHD Testbed

Type	Make	Model	Radio	Spatial Stream	Quantity
Smartphone	Samsung	Galaxy S4	Broadcom 4335	1SS	100
Ultramobile	Apple	MacBook Air	Broadcom 4360	2SS	100
Laptop	Apple	MacBook Pro	Broadcom 43460	3SS	100

You clearly can see that the 10 station throughput on the left is very similar to what we computed in the previous section. Some degradation is visible because we are beginning with 10 concurrent devices instead of a single station. Performance degrades further the more clients are added, by an average of 20% at 50 clients, and up to 60% at 100 clients.



To read a detailed discussion on the Aruba VHD testbed and see a photo of what it looks like, see [Appendix T-A: Aruba Very High-Density Testbed](#) of the *Very High-Density 802.11ac Networks Theory Guide*.

This performance drop occurs with equipment from all WLAN manufacturers; it is a basic feature of Wi-Fi operation at high densities. You will see similar results from other companies.

Using Lab Results for VHD Capacity Planning

You can apply these lab results directly to estimate system capacity in your own VHD environments.

The middle term in the TST formula is Average Channel Throughput. You can obtain this from the lab results in [Figure EC2-2 on page 16](#) by applying an adjustment to reflect real-world interference and other RF challenges.

$$\text{Average Channel Throughput} = \text{Lab Throughput (N)} * (1 - \text{Impairment Factor}\%) \quad (4)$$

Where:

- Lab Throughput = The measured multiclient throughput in a clean lab environment for N concurrent devices
- Impairment Factor = Percentage reduction in lab throughput expected for specific VHD facility type

Aruba recommends you use different impairment values for the 2.4-GHz band and 5-GHz band. ([Table EC2-12 on page 27](#) lists suggested impairment factors based on Aruba experience.)

To simplify the process, let us convert the right side of [Figure EC2-2](#) into a more useful tabular format in [Table EC2-5](#). Values have been rounded to the nearest 5 Mbps for convenience.

Table EC2-5 Multiclient Lab Throughput by Spatial Stream Mix

Spatial Stream Mix	50 Concurrent	75 Concurrent	100 Concurrent
100% 1SS Device	45 Mbps	35 Mbps	30 Mbps
75% 1SS + 25% 2SS	60 Mbps	40 Mbps	30 Mbps
50% 1SS + 50% 2SS	70 Mbps	50 Mbps	40 Mbps
30% 1SS + 60% 2SS + 10% 3SS	75 Mbps	50 Mbps	40 Mbps
100% 2SS Device	100 Mbps	65 Mbps	40 Mbps
100% 3SS Device	160 Mbps	115 Mbps	80 Mbps

You will see Table EC2-5 many times in the VRD. To use this table for capacity planning via the TST process, first choose a row based on the spatial stream mix you expect for your facility. Then choose a column based on the number of concurrent devices you expect to *attempt to use the channel at the same time*. This number is determined easily from the device duty cycle you anticipate.

We will condense this table still further in step 2 of the TST process. But first we must review two more building blocks: how to estimate the spatial stream mix, and how to calculate duty cycle.

Modeling the Spatial Stream Mix of VHD Areas

The preceding sections show the wide variation in potential single-client and multiclient performance from devices of different spatial stream capabilities. To complicate matters further, real wireless networks have a wide variety of devices that use them. You cannot tell how many streams a device supports based on the device type. Smartphones and tablets both come in 1SS and 2SS flavors now. Laptops generally come in 2SS and 3SS flavors. How can the wireless engineer forecast the device mix so that an accurate capacity plan can be generated?

Fortunately, we do see common trends in device form factors for specific types of VHD facilities. We can use these trends to abstract a couple of spatial stream mixes that meet our planning needs.

Table EC2-6 Device Type Distribution in Common VHD Facility Types

VHD Venue Type	Phones	Tablets	Laptops
Classroom / Lecture Hall	40%	30%	30%
Convention Center	40%	30%	30%
Airport	40%	30%	30%
Casino	90%	10%	-
Stadium / Arena /Theater	90%	10%	-
Shopping Mall	90%	10%	-

Table EC2-7 Spatial Stream Mix Forecast for Common Mobile Devices

Spatial Streams	1SS	2SS	3SS
Smartphone	50%	50%	-
Tablet	30%	70%	-
Laptop	-	70%	30%
Wearable	100%	-	-

[Table EC2-6](#) shows estimated breakdowns of device types for six major types of VHD facility. These forecasts represent a consensus estimate of the Aruba CTO and product management teams because precise market research is not available. As you look at the table, remember that the average number of devices a person has varies in each facility type.

No clear relationship exists between a device type and the number of spatial streams it supports. [Table EC2-7](#) is one scenario for how the mix of device spatial streams may evolve over the next few years. Currently, industry analysts expect that most client devices including smartphones will become 2SS over time. This increase in 2SS devices is due to 802.11ac MU-MIMO being optimized for four clients of 2SS each. We include the wearable category because these devices will be a factor in the future. To the extent that they are Wi-Fi-enabled instead of Bluetooth, wearable devices are certain to be 1SS for battery and space reasons. As this VRD went to press, Apple announced that the iWatch supports Wi-Fi operation.

Table EC2-6 and Table EC2-7 may be combined to produce a blended spatial stream mix forecast for each of the six VHD facility types. Table EC2-8 is the result. It neatly boils down all of the complexity of the other tables into a simple, easy-to-use resource. This is the critical table, and you will use it in every VHD capacity plan you prepare.


Table EC2-8 Spatial Stream Blend Lookup Table

VHD Venue Type	Application Usage Profile	Devices/ Person (Now)	Devices/ Person (Future)	1SS %	2SS %	3SS %
Classroom / Lecture Hall	Work / Study	3	5	30%	60%	10%
Convention Center						
Airport						
Casino	Fan / Guest	1	2	50%	50%	--
Stadium / Arena / Theater						
Shopping Mall						

We use these two blends repeatedly in examples in this chapter. Of course, this model may not apply in every country and the variation between regions could be significant. You may want to model the blend for your own country and end-user profiles once you become comfortable with the overall approach.

Estimation Process for System Throughput

Now that terminology is defined and the basic capacity concepts are explained, we work through the TST estimation process in detail from start to finish. Using the structure of formula (1), we follow this five step approach:



Step 1 – Select Channel Count

The goal of this step is to determine how many channels the VHD network will use. Using more channels directly increases capacity.

To keep this chapter simple and focused, it is assumed that:

- The wireless architect deploys 20-MHz channels.
- DFS channels are employed if they are available (unless spectrum surveys show significant radar events).
- 5-GHz is used as the primary service band.

The next chapter provides an in-depth exploration of 20-MHz vs. 40-MHz vs. 80-MHz channels, and it describes the specific reasons why bonded channels should not be used in VHD areas. For now, we treat this as an assumption.

Stacking Channels Increases Capacity

In any VHD WLAN, we must use as many radio channels as possible because capacity increases linearly with the number of channels. [Figure EC2-3](#) shows that two co-located APs on different channels provide roughly twice the capacity of a single AP. With three APs on different channels in the same RF collision domain, capacity is roughly tripled, and so on.




Figure EC2-3 Using Additional Channels to Increase WLAN Capacity

Therefore, the number of channels is the primary capacity constraint on a VHD WLAN. For this reason, VHD WLANs should always use the 5-GHz band for primary client service because most countries have many more channels in this band. [Figure EC2-4](#) shows the number of 80-MHz, 40-MHz, and 20-MHz channels supported by 802.11 for use in the 5-GHz band as of the date of publication.




Figure EC2-4 Allowed Channels In US Regulatory Domain

Choosing the Number of Channels

In short, your VHD network should use every single 20-MHz channel allowed in your country including DFS channels. This rule has only a few exceptions.

The usable channel count is determined as follows:

$$\text{Usable Channels} = \text{Allowed Channels} - \text{Impaired DFS Channels} - \text{Ch. 144} - \text{Reserved Channels} - \text{Other Impaired Channels}$$

Where:

- Allowed Channels = Total number of channels permitted by the spectrum regulator in a given country
- Impaired DFS Channels = The number of DFS channels on which radar has been proven to exist
- Channel 144 = A single channel that should not be used at this time
- Reserved Channels = Private “house” channels for the VHD facility
- Other Impaired Channels = Unusable channels due to recurring interference proven with a spectrum analyzer, or due to device incompatibilities

The number of available 5-GHz channels varies significantly from country to country. [Appendix EC-A: Worldwide 5-GHz Channel Availability as of March 1, 2015](#) lists the typical channels available for some common countries and regions at the time of publication.

For countries or regions that are not listed, contact the Aruba Technical Assistance Center or a professional installer. An Aruba controller also reports the valid channels for a given regulatory domain with the “show ap allowed-channels country-code <country code> ap-type <AP model>” command.

[Table EC2-9](#) lists the total number of non-DFS and DFS channels in some common regions, as of the date of publication.

Table EC2-9 Available 5-GHz Channel Counts in Selected Countries or Regions

Channel	United States	Brazil	Europe & Turkey	South Africa	China	Japan	Korea	Singapore	Taiwan	Australia	New Zealand
TOTAL NON-DFS	9	9	4	4	9	4	8	9	9	9	9
TOTAL DFS	12 ¹ / 15	15	15	15	4	15	12	15	11	13	15
TOTAL	21 ¹ / 24	24	19	19	13	19	20	24	20	25	24

1. These channels were temporarily disallowed in 2013-2014 in the US. APs released from 2015 on may use these channels if they pass DFS certification.

Over the last few years the available Wi-Fi channels in most countries has increased significantly. Stay current on the latest changes in the regions you support.

Exception #1 – DFS Channels

DFS channels should almost always be used in VHD areas. In general, the capacity benefit of using these channels far outweighs the potential costs. But this decision must be made on a per-facility basis.

Channels subject to DFS rules vary slightly from country to country, but they usually include 52 – 64 and 100 - 144. These channels are called DFS channels because special rules called “Dynamic Frequency Selection” must be enforced.



For an explanation of Dynamic Frequency Selection (DFS), see [Appendix EC-C: DFS Surveys and Operating Rules](#).

In the past, there were good reasons not to use DFS channels. However, significant developments have changed our guidance. Aruba has large VHD deployments that use DFS channels successfully.

You must apply three basic criteria when you decide whether or not to use DFS channels: client compatibility, voice roaming, and radar exposure. This topic is of interest to engineers and architects, so we explore it fully in this guide.

Client Capabilities

The first and most critical concern you must evaluate to use DFS channels is the capabilities of your expected client devices. As a general matter, as of 2015, the vast majority of newly introduced mobile devices support DFS channels. All five major mobile operating systems in the world now support DFS (Table EC2-10).

Table EC2-10 DFS Channel Support by Operating System in 2015

Platform	Operating System	DFS Support	Comment
Phone / Tablet	Apple iOS	Yes	iOS has supported DFS channels for several years.
	Google Android	Yes **	Most if not all new 802.11ac Android devices support DFS because it is required for 80-MHz channels. ** Most legacy 802.11n Android devices do not.
	Microsoft Windows Phone	Yes	
Laptop / Netbook	Apple MacOS	Yes	MacOS has supported DFS channels for several years.
	Microsoft Windows	Yes	Windows has supported DFS channels for several years.

The addition of DFS support to 802.11ac Android devices in 2014 removed one of the last barriers to using DFS channels. While some low-cost manufacturers of Android devices have not enabled these channels, Android itself fully supports DFS operation.

Ultimately, this decision comes down to what part of the world you are in. In the US and Western Europe, Apple iOS devices are the dominant operating system in most high-density areas. Devices also tend to refresh more quickly. In these countries, it is generally safe to adopt DFS channels unless you are certain you have a radar exposure.

In Asia, Eastern Europe, the Middle East, and Africa, older Android devices tend to dominate. Refreshes are much more gradual, so many of your potential users may not be able to see the DFS channels. This will change over the next 2 to 3 years.

Voice Roaming

VHD environments almost never require voice roaming. They generally are designed for stationary users who are sitting or standing.

If you require voice roaming in your VHD environment, do not use DFS channels due to longer handover times.

Roaming by users who are streaming video generally works fine with DFS channels because the video buffers exceed the handover times.

Radar Exposure

The third concern about using DFS channels is whether the location of your VHD network is vulnerable to actual radar events on some or all DFS channels. The risk is that a serving AP could be temporarily “knocked out” during a public event for 60-90 seconds while it switches channels.

This presence or absence of radar can be easily checked by conducting a DFS survey. The DFS survey procedure is explained in [Appendix EC-C: DFS Surveys and Operating Rules](#).

Do not assume that simply because your VHD facility is near an airport, military base, or a body of water with shipping traffic that DFS channels are not usable. There may be no radar installation at all, or radar may be present only on specific frequencies, which leaves all other frequencies available for use.

You may also find that only certain parts of your facility experience radar events. In an outdoor stadium for example, the upper seating levels may be more likely than lower bowl seating to experience radar from a nearby source. An indoor venue may see radar events outside but not inside due to attenuation of signal through the building walls.

A DFS survey tells you which channels, if any, should be excluded from your channel plan. Exclude only those channels that experience daily, continuous or recurring radar events. Infrequent radar events do not justify ruling out the use of those channels.

After the network has been deployed, future radar events show up in the ARM history and the system log. These should be periodically monitored for changes, and it’s a good idea to set up a SYSLOG alert specific to radar messages. This should provide you with peace of mind that if an event does occur, that you will know about it.

How Do You Balance the Risks and Rewards?

Basically, the risks of DFS are client compatibility and channel interruption. Of the two risks, client compatibility is the most serious. If you know for certain that a large percentage of your device population cannot see DFS channels, then it is foolish to use those channels because you will have ongoing complaints. If you decide against DFS for this reason, monitor your population over time and revisit the idea in future years.

Channel interruption is only a factor if the network experiences verified, recurring, frequent radar events. In this case, it is wise to exclude those channels from your plan.

However, with a clean DFS survey, the risk of channel interruption is very low. Imagine that a few of your APs happen to experience a radar event when the facility is full of users. The radar event affects only one or two channels. By definition, many more APs cover the area because you have stacked channels. So clients have many other choices of APs. In this case, the good of the many greatly outweighs the good of the few.

It is a worthwhile price to pay to occasionally have a small number of clients disconnected momentarily in order to reduce the overall client load on the non-DFS channels. When you avoid DFS channels, you ensure that every one of your users gets less airtime, all the time. When you use DFS channels, you can cut the average per-channel device load by half or more (depending on the particular country). This reduced load increases total system capacity and performance, all the time.

Considering client compatibility, it is likewise worthwhile to use DFS channels even if you know that as many as 10-15% of your clients cannot use those channels. As long as you ensure a regular mix of both non-DFS and DFS channels in your channel plan, those devices will be able to connect and use the system. Perhaps they will experience a slightly slower connect rate.

Ultimately, VHD networks are about tradeoffs. Cost, capacity, aesthetics, and schedule must be considered. Similarly, performance and compatibility must be weighed. All VHD networks are imperfect by definition due to the huge oversubscription they must be designed to support. The question is how to

deliver the maximum experience to the largest number of users possible. Viewed this way, the bias should be toward using DFS channels.

Exception #2 – Channel 144

In general, do not use channel 144 for guest access until the client mix in your particular venue is over 50% 802.11ac capable. Most venues in the US will not reach this client mix until sometime in 2016. For the rest of the world, this may take another one or two years.

The reason is that while most 802.11n and nearly all 802.11ac devices on the market support DFS channels, no 802.11n device can see channel 144. Channel 144 support was added as part of the 802.11ac amendment.

One good use of channel 144 in the meantime is as a “house” channel if you can ensure that all devices that use that channel support it.

Exception #3 – “House” Channels

The concept of reserving one or two “house” channels is discussed in various chapters of this VRD. You might reserve house channels if you expect to support regular presentations in a large convention center that depend on Wi-Fi and cannot be impacted by guests. Stadiums and arenas sometimes do this to guarantee a clean channel for use by team applications that run on mobile devices.

In general, Aruba recommends against reserving house channels unless certain house users consistently have trouble. When a channel is removed from service, the remaining channels must carry more load. With a good DFS design, you might find that the average per-channel load is low enough to support guest and house applications.



To reserve house channels, you must use DFS channels. Without DFS, you do not have enough channels to serve guests.

House channels should be deducted from the channel count used for the TST calculation.

Exception #4 – High Interference Impairments

Because VHD coverage zones are almost always part of larger facilities, the channel plan for the rest of the site may also impose constraints on channel availability. Be sure to consider any reserved channels that are required for indoor or outdoor mesh operations or for high-duty-cycle applications such as wireless IP video.

Before you commit to a final channel plan, you must use a professional grade spectrum analyzer and a packet capture utility to verify the condition of the air in the area to be covered. The spectrum analyzer shows the energy from all transmitters in the band – both Wi-Fi and otherwise. The packet capture utility shows the level of Wi-Fi specific load on each channel.

Aruba recommends that you analyze the air when the venue is empty and during several live events. It is common to find other non-Wi-Fi transmitters that use the same unlicensed frequencies during events. Coach communication systems, wireless microphones, wireless mice and keyboards, and aerial video cameras are just some of the interference sources we have encountered that appear only during events.

The key here is to understand the duty cycle of the interference sources, and the frequencies that are affected. Some interference is tolerable if it is infrequent. When high-duty-cycle sources are discovered, either those channels must be taken out of service, or the interferer must be disabled permanently, or an alternate technology for the interferer to use must be found that does not overlap with the Wi-Fi system.

Exception #5 – Interoperability Impairment

Certain older point-of-sale terminals and voice handsets do not support channel 165 in some countries. Verify that the equipment you plan to support can use this channel. If not, you may need to remove it from the channel plan.

The 2.4-GHz Band

The entire discussion in this chapter so far has dealt with the 5-GHz band. Aruba recommends that you plan for 5 GHz as the primary service band.

It is assumed that guest service will be provided on the 2.4-GHz band in any VHD area. However, the 2.4-GHz band primarily is for users with older legacy devices are incapable of 5-GHz operation.

Technically, using the 2.4-GHz band increases your available channel count by three in countries that support 11 channels in 2.4-GHz, and by four in countries that allow 13 channels.

However, the 2.4-GHz band is increasingly unusable in VHD areas due to Wi-Fi and non-Wi-Fi interference, especially Bluetooth. The more seats in the facility, the worse the problem. [Figure EC2-5](#) is a 10-minute spectrum capture from a 50,000 seat stadium. The upper left quadrant shows the 2.4-GHz band, while the other quadrants show the various 5-GHz sub-bands.




Figure EC2-5 Spectrum Analysis of 2.4-GHz and 5-GHz Bands in a Large Stadium

In this capture, the noise floor has effectively been raised to -80 dBm. The main cause is the very large number of Wi-Fi APs using 802.11b modulations, which results in a nearly continuous duty cycle of transmissions across the 2.4-GHz band (the solid line across the center). The 802.11b CCK waveform has a distinctive “scalloped” edge, as compared with 802.11n OFDM waveform which has a much flatter “top” and steep sides.

If you look carefully in the plot in [Figure EC2-5](#), you can actually see multiple “tiers” of 802.11b APs on channels 1, 6, and 11. These are the nested scalloped patterns at different signal levels. Each of these signal groups is a batch of 802.11b APs at different distances from the point of measurement.

As a result, we plan capacity separately for 5 GHz and 2.4 GHz. We choose a different average channel throughput value for each band. You'll learn how in the next section.

Step 2 – Estimate Unimpaired Multi-client Throughput

The purpose of this step is to estimate the average usable capacity of a Wi-Fi channel under load in clean conditions. You use the Aruba VHD lab test data presented earlier in the chapter. This value is multiplied by the channel count to determine the TST.

Table EC2-11 contains all the information you need to complete this step. The unimpaired goodput values in the far right column come directly from the Aruba VHD lab test results presented in Figure EC2-2 on page 16. For each of the six VHD facility types discussed so far, we have selected both a particular curve from the figure as well as a single point on that curve. The curve was chosen based on the spatial stream mix expected in that environment listed in Figure EC2-8 on page 35. The exact throughput value selected on the curve was chosen based on our estimate of the typical number of concurrent users per channel in those six environments. Use this information as a guide; as the wireless engineer, you must use your professional judgment to adapt the table for your particular venue.

Table EC2-11 Unimpaired Channel Throughput Reference Table

VHD Venue Type	Application Usage Profile	Spatial Stream Mix	Typical Concurrent Users per Channel	Unimpaired VHT20 Goodput
Classroom / Lecture Hall	Work / Study	1SS – 30%	50	75 Mbps
Convention Center		2SS – 60%	75	50 Mbps
Airport		3SS – 10%	75	50 Mbps
Casino	Fan / Guest	1SS – 50%	50	70 Mbps
Stadium / Arena		2SS – 50%	100	40 Mbps
Shopping Mall			50	70 Mbps

The column on the right is the number you should choose for this step, adjusted for your facility. Remember, we made these assumptions earlier:

- Quality RF design with consistent SINR ≥ 25 dB through the VHD area
- Maximum MCS data rates are possible over 90% of the VHD area
- 802.11ac with VHT20 modulations are being used

If any of these assumptions are not true in your venue, use your judgment to make further adjustments. The result should be a per-channel average throughput value that instinctively feels “right” for the specific scenario you are engineering.

Step 3 – Choose and Apply Impairment Factor

The throughput values in Table EC2-11 are from a clean lab environment, without any of the interference or other impairments that are so common in VHD areas. So these numbers cannot be used directly in the TST formula. They must be adapted based on the characteristics of each particular facility.

To adapt the unimpaired throughput values, we need a metric to capture the effect of CCI, ACI, non-Wi-Fi interference, and a safety margin for the unexpected, such as more devices than planned or heavier transmit duty cycles. The function is similar to the fade margin in an RF link budget. Aruba calls this metric the “impairment factor”.

You should choose a different impairment for 2.4-GHz and 5-GHz bands, due to the significantly higher interference level in the 2.4-GHz band worldwide. We suggest you use a flat figure based on the type of facility being covered. [Table EC2-12](#) includes consensus values from the Aruba high-density engineering team based on our deployment experience.

Table EC2-12 Suggested Impairment Reference Table

VHD Venue Type	Suggested 2.4-GHz Impairment	Suggested 5-GHz Impairment	Rationale
Classroom / Lecture Hall	10%	5%	• Above average duty cycles • Little or no reuse of channels in the same room • Structural isolation of same-channel BSS in adjacent rooms • Minimal My-Fi usage
Convention Center	25%	10%	• Moderate duty cycles • Significant numbers of same-channel APs • Large open areas with direct exposure to interference sources • Non-Wi-Fi interferers • Higher My-Fi usage in booth displays, presenters, attendees
Airport	25%	15%	• Lower duty cycles (except for people streaming videos) • Structural isolation of same-channel BSS in adjacent rooms • Heavy My-Fi usage
Casino	25%	10%	• Low duty cycles on casino floor • Low My-Fi usage
Stadium / Arena	50%	25%	• Low-to-moderate duty cycles • Significant numbers of same-channel APs • Large open areas with direct exposure to interference sources • Non-Wi-Fi interferers • High My-Fi usage
Shopping Mall	10%	5%	• Moderate duty cycles • Significant numbers of same-channel APs • Large open areas with direct exposure to interference sources • Non-Wi-Fi interferers

The next step is to apply the impairment value to the unimpaired throughput value from step 2. This produces the values that you will plug into the TST formula.

Table EC2-13 Impaired Channel Throughput Reference Table

VHD Venue Type	Unimpaired VHT20 Goodput	Suggested Impairment		Impaired Goodput	
		2.4 GHz	5 GHz	2.4 GHz HT	5 GHz VHT
Classroom / Lecture Hall	75 Mbps	10%	5%	67 Mbps	71 Mbps
Convention Center	50 Mbps	25%	10%	37 Mbps	45 Mbps
Airport	50 Mbps	25%	15%	37 Mbps	42 Mbps
Casino	70 Mbps	25%	10%	52 Mbps	63 Mbps
Stadium / Arena	40 Mbps	50%	25%	20 Mbps	30 Mbps
Shopping Mall	70 Mbps	10%	5%	63 Mbps	66 Mbps

This table is simply meant to be an example to show you how to calculate the impaired values yourself. Your environment may be better or worse than these impairment estimates. Obtain spectrum and packet captures as part of your design process. Analyze this data and use your best judgment in your own circumstances. If your facility has especially light interference, you can use a smaller impairment figure. Use a larger figure if the air is particularly dirty.

The expected concurrent users in your environment may also differ from the example in step 2. For example, for the convention center type, we have selected 75 concurrent users per channel. For a very large hall over 5,000 m² (54,000 ft²), it would be appropriate to go with 100 users per channel instead.

The bottom line is that you must adapt the entire methodology that is shared in this chapter for each and every VHD facility you design. Aruba can offer no one-size-fits-all table.

Step 4 – Select Spatial Reuse Factor

The next important variable in the total system throughput equation is the spatial reuse factor. Because spatial reuse is so fundamental to VHD networks, and so widely misunderstood, we provide a high-level introduction to the topic here. For a far more detailed explanation including the mechanics of CCI, WLAN architects should also read [Chapter T-5: Understanding RF Collision Domains](#) of the *Very High-Density 802.11ac Networks Theory Guide*.

What is RF Spatial Reuse?

Wireless signal strength decays over distance, so a given radio channel can be reused at intervals. This concept has long been used by mobile telephone networks, and it is central to most WLAN architectures. All enterprise WLANs reuse channels in clusters to serve large areas where the radios are separated from one another by free space, walls, or other structures. In this case, the purpose of reuse is to provide a consistent signal level everywhere in a facility, regardless of the actual number of client devices. [Figure EC2-6](#) shows two channel reuse clusters and the relative position of reused channels.




Figure EC2-6 Channel Plan with 11 Channels in 5 GHz with Minimum Separation of Two Cells

In this scenario, two devices on the same channel but in different cells can transmit at the same time. In this case, “RF spatial reuse” is achieved because the radio signals of one cell do not interfere with the other cell that is some distance away.

However, in the VHD environment we are confronted by relatively small physical areas (in radio terms) that have many APs on the same channel to achieve the ADC target. In this case, the entire VHD area may be thought of as a single cell whose capacity must be shared.




Figure EC2-7 Stacking Same-Channel APs in a Large Indoor Arena

In [Figure EC2-7](#), we have an indoor arena with four APs on channel 48, all of which can hear the others with high enough SINR to decode their transmissions. Due to the closed roof, signals bounce well inside, which makes the entire arena a single RF collision domain.

In this case, even though we are reusing the channel *number*, we are (in most cases) not reusing the radio *spectrum*.

This concept is important because it leads to a common misunderstanding that squeezing in more APs somehow increases the capacity of a VHD network. It is true that having more APs does increase the number of associated devices the network can absorb. However, the TST is not changed if the APs, their clients, or both can hear one another.

Defining the Reuse Factor

In simple terms, the reuse factor is the number of devices that can use the same channel at exactly the same time, without interfering with one another.

For example, in a typical 20,000 seat arena it is common to find 100 APs in the bowl seating area. If 20 channels are being used in 5 GHz, then on average every channel has five APs. If those APs **and their associated clients** are completely RF-isolated from one another, then the reuse factor equals 5.

However, in practice, reuse is extremely difficult if not impossible to achieve in indoor VHD areas. More often than not, you should choose a reuse factor of 1. (It is actually possible to be less than 1 in certain co-channel interference conditions.)

Outdoor VHD areas that are open to the sky can sometimes achieve a reuse factor of more than 1. For example, let us imagine a soccer stadium with four devices on the same channel, one in each corner of the seating bowl. If all four devices can transmit simultaneously, then the reuse factor is 4.

Step 5 – Calculate Total System Throughput

You have determined the channel count, the average impaired per-channel throughput, and the reuse factor. It is now possible to compute the system throughput for the VHD area using formula (1).

We will work through three specific example scenarios momentarily. But first, let us examine how to convert the TST into a per-device value.

Estimation Process for Per-Device Throughput

As stated at the top of this chapter, most customers who purchase a VHD system define their requirements in terms of minimum per-seat or per-device throughput. When video or other high-bitrate services are required, some customers even attempt to guarantee such minimums contractually.

Aruba's top-down capacity planning method must be translated into a good-faith estimate of per-device performance, which is provided by formula (2):

$$\text{Average Device Throughput} = \frac{\text{Total System Throughput}}{\text{Instantaneous User Count}}$$

The formula shows that average device throughput (ADT) is equal to the TST divided by the instantaneous user count. Instantaneous user count is the percentage of total associated devices that attempt to use the system at any given moment. Basically, instantaneous user count is the duty cycle of the device transmissions. So we can rewrite the formula like this:

$$\text{Average Device Throughput} = \frac{\text{Total System Throughput}}{\text{Active Device Capacity} * \text{Device Duty Cycle}}$$

You already have computed the guest ADC, staff ADC, and total ADC during the dimensioning step in [Chapter P-2: System Dimensioning](#) of the *Very High-Density 802.11ac Networks Planning Guide*.

Device duty cycle is simply the percentage of time that a given device is attempting to transmit. This time includes when management and control traffic and data frames are sent, and retries of any of these that fail to get through.

Unfortunately, estimating duty cycle in WLANs is somewhat subjective for several reasons. The first reason is that no two VHD areas have exactly the same user behavior. Students in lecture halls in the US behave differently than those in Asia or Europe. Football fans behave differently than basketball or hockey fans. Music concert behavior depends on the age of the audience.

Second, everyone has a different mix of applications on their devices and subscribes to different online services. At any given time, a user's Wi-Fi device may be running dozens of applications, each one with a different network usage profile.

A third reason is that as the number of Wi-Fi-enabled devices per person increases, only one device at a time will be “primary” in the user’s hands. Other devices in a pocket, purse, or backpack may use the network, but in a less aggressive manner.

Still, the wireless engineer must choose some duty cycle value in order to make a capacity plan. [Table EC2-14](#) lists the five duty cycles that Aruba has abstracted for VHD areas based on the consensus estimate of our VHD engineering team.

Table EC2-14 Estimated Duty Cycles for VHD Areas

Category	Duty Cycle	User & Device Behavior Examples	Usage Mode
Background	5%	Network keepalive / App phonehome	Secondary
Checking In	10%	Web browsing / Checking email / Social updates	Primary
Semi-Focused	25%	Streaming scores / Courseware / Online exam	Primary
Working	50%	Virtual desktop / HTTPS application / Terminal	Primary
Active	100%	Video streaming / Voice streaming / Gaming	Primary

You must use your own judgment to apply this table to your own facilities. This information is meant to be a starting point to show you the way. At the end of the chapter, these duty cycle values are used in three specific examples.

So for example, consider a 5,000 seat convention center ballroom with a TST of 500 Mbps, a 50% take rate and a 25% duty cycle.

$$ADT = \frac{500 \text{ Mbps}}{5,000 \text{ Users} * 50\% \text{ Take Rate} * 25\% \text{ Duty Cycle}} = \frac{500}{(5,000 * .5 * .25)} = 800 \text{ Kbps}$$

The vast majority of the work for the VHD capacity plan is calculating the TST. After that is complete, you can produce a per-device estimate quickly.

Examples

You have learned how to estimate system throughput, per-channel throughput, and per-device throughput for a VHD system. Now we work through three different examples to highlight how the process works in large and small scenarios.

Example #1 – Small Auditorium

In a small auditorium with 500 seats, it is not necessary to use very many APs. Therefore, we do not need many channels. We assume:

- Average of three devices per person
- 70% take rate
- 50% / 50% split between 5-GHz and 2.4-GHz band
- 150 associations per radio

Using the VHD dimensioning process from [Chapter P-2: System Dimensioning](#) of the *Very High-Density 802.11ac Networks Planning Guide*, yields this ADC value:

$$\text{ADC} = 500 \text{ Seats} * 3 \text{ Devices} * 70\% \text{ Take Rate} = 1,050 \text{ Devices}$$

These devices are split evenly between the two frequency bands, with 525 devices on each one. We can obtain the required radio count by this equation:

$$\text{APs} = 5\text{-GHz Radio Count} * 525 \text{ Devices} \div 150 \text{ Associations per radio} = 4$$

So far, so good. Therefore, four dual-radio APs are enough from an association capacity perspective. Obviously, we need not worry about whether or not to use DFS channels in this example.

Now let us calculate the total system throughput. To do that, we must make a few more assumptions:

- Each user has one smartphone, one tablet, and one laptop
- Unimpaired throughput of 75 Mbps using Lecture Hall spatial stream mix per [Table EC2-11 on page 26](#). (30% / 60% / 10%)
- Impairment factor of 5% for 5 GHz and 10% for 2.4 GHz (from [Table EC2-13 on page 27](#))
- The classroom learning applications have a duty cycle of 25% (from [Table EC2-14 on page 31](#))

When we follow the process laid out in this chapter, we begin to estimate the TST by determining the channel count. This count must be four for 5 GHz, however the 2.4-GHz band has only three usable channels. So we begin to fill out a throughput table like [Table EC2-15](#).

Table EC2-15 System and Device Throughput for Auditorium Example

Band	ADC	Radios	Channels	Channel Throughput	Reuse Factor	Total System Throughput	Duty Cycle	Per-Device Throughput
5 GHz	525	4	4	71 Mbps	1	284 Mbps	25%	2.1 Mbps
2.4 GHz	525	4	3	67 Mbps	1	201 Mbps	25%	1.9 Mbps
			7			485 Mbps		

As you can see, the 5-GHz band has a total of 284 Mbps of possible system throughput, and the 2.4-GHz band has another 201 Mbps. These numbers are cumulative, so the potential load that the system can generate on the WAN uplink is nearly 500 Mbps.

We can also obtain an estimated per-device throughput by using the duty cycle value. A 25% duty cycle means that on average, one-fourth of the devices attempt to use the system at any one time. So we divide 25% of the total ADC into the TST.

Example #2 – Indoor Arena in China

This example is designed to explore channel reuse and per-country limits on channel availability.

Consider an indoor arena with these attributes:

- 10,000 seats
- 50% take rate
- One device per person
- 75% / 25% split between 5-GHz and 2.4-GHz bands

- Unimpaired throughput of 40 Mbps using Arena spatial stream mix per [Table EC2-11 on page 26](#) (50% / 50% / 0%)
- Impairment factor of 25% for 5-GHz and 50% for 2.4-GHz (from [Table EC2-13 on page 27](#))
- User devices have a duty cycle of 10% (from [Table EC2-14 on page 31](#))
- Indoor channel model with closed roof (no RF spatial reuse)

Therefore, we must support a total associated device count of $10,000 * 50\% * 1 = 5,000$. Of these, three-quarters are on 5-GHz band and one-quarter are on 2.4-GHz band.

In Example #1, we did not need to consult the country table because we knew in advance that the room did not require many channels. But in this example, it is already clear that we are going to have more APs than channels to meet the associated device capacity target.

Consulting [Appendix EC-A: Worldwide 5-GHz Channel Availability as of March 1, 2015](#), it seems that China currently supports 13 indoor channels on 5-GHz.



Be sure to always check for the updated regulatory information for your country. New channels are coming online for Wi-Fi use around the world. For instance, China only allowed 5 channels until 2013.

Turning to average bandwidth, [Table EC2-13 on page 27](#) gives us a value of 30 Mbps for the Arena case in 5-GHz band, and 20 Mbps for 2.4-GHz band.

The reuse factor should be set to 1. No spatial reuse is possible in most indoor venues, even those as large as 20,000 seats!

So the throughput table looks like this:

Band	ADC	Radios	Channels	Channel Throughput	Reuse Factor	Total System Throughput	Duty Cycle	Per-Device Throughput
5 GHz	3,750	25	13	30 Mbps	1	390 Mbps	10%	1.0 Mbps
2.4 GHz	1,250	9	3	20 Mbps	1	60 Mbps	10%	480 Kbps
			16			450 Mbps		

Therefore, we have a total estimated system bandwidth of just 450 Mbps for the entire venue. Note that each channel number in the 5-GHz band is reused twice, and each 2.4-GHz channel number is reused three times.

Example #3 – Outdoor Stadium in USA

The United States generally leads the way for setting aside new spectrum for Wi-Fi. More channels are available here than almost anywhere else. In this example, we explore capacity with a maximum channel allocation and different reuse factors.

Consider an outdoor stadium with these attributes:

- 60,000 seats
- 50% take rate
- One device per person
- 75% / 25% split between 5-GHz and 2.4-GHz bands

- Unimpaired throughput of 40 Mbps using Stadium spatial stream mix per [Table EC2-11 on page 26](#) (50% / 50% / 0%)
- Impairment factor of 25% for 5-GHz and 50% for 2.4-GHz bands (from [Table EC2-13 on page 27](#))
- User devices have a duty cycle of 10% (from [Table EC2-14 on page 31](#))
- Outdoor channel model with open roof

Therefore, we must support a total ADC of $60,000 * 50\% * 1 = 30,000$ devices. Of these, three-quarters are on 5-GHz band and one-quarter are on 2.4-GHz band.

From [Appendix EC-A: Worldwide 5-GHz Channel Availability as of March 1, 2015](#), we see that the US presently allows 21 20-MHz channels in the 5-GHz band (excluding channels 120-128 and 144). From [Table EC2-13 on page 27](#), we choose the same average channel throughput value as the previous example. Channel usage depends on usage model, not geography.

For reuse factor, due to the outdoor scenario with no roof, assume that it may be possible to achieve some spatial reuse. For example, two users on opposite ends of the field may be able to transmit at the same time on the same channel. So let us use a reuse factor of 2.

Band	ADC	Radios	Channels	Channel Throughput	Reuse Factor	Total System Throughput	Duty Cycle	Per-Device Throughput
5 GHz	22,500	150	21	30 Mbps	2	1,260 Mbps	10%	560 Kbps
2.4 GHz	7,500	50	3	20 Mbps	2	120 Mbps	10%	160 Kbps
			24			1,380 Mbps		

Depending on the exact RF coverage strategy chosen, the reuse factor may be higher or lower. Aruba research suggests that the reuse factor in an outdoor venue can range from 2 for some overhead designs to more than 8 for a properly designed underseat picocell. By varying the reuse factor variable, we can perform a “what if” analysis of a range of capacities:

Reuse Factor	1.0	2.0	3.0	4.0	5.0	6.0
Stadium - 21 channels	630 Mbps	1,260 Mbps	1,890 Mbps	2,520 Mbps	3,150 Mbps	3,780 Mbps

So when planning for applications and uplink bandwidth, you must understand exactly how the RF design will perform. The minimum WAN uplink in this scenario is 1 Gbps for reuse factor of 1, and up to 4 Gbps for a reuse factor of 6.

At a per-device level, the plan suggests that the VHD system can deliver an **average** of 560 Kbps per device on 5-GHz, and just 160 Kbps on the 2.4-GHz band. Of course, this example has dramatically more users than the last example. So to deliver a usable experience, it is of the highest importance that the duty cycles in these very large venues be comparatively low. Here are the per-device throughputs by band for a range of other duty cycle values.

Duty Cycle	5%	10%	15%	20%
5 GHz	1,120 Kbps	560 Kbps	420 Kbps	280 Kbps
2.4 GHz	320 Kbps	160 Kbps	120 Kbps	80 Kbps

These numbers are astonishingly low. So if the 10% assumption is not correct, it will have a huge effect on these figures. Also, if the final system fails to deliver the planned RF spectrum reuse, these figures will likewise be severely reduced.

Such numbers have enormous consequences for capacity planning for in-stadium fan applications that want to deliver video replays as well as other in-stadium services.

Comparing Examples to Aruba VHD Lab Results

These three examples demonstrate the huge potential difference in total system throughput and per-device throughput for VHD facilities of various kinds and sizes. Some customers may have difficulty accepting how low the per-device numbers get in the larger facilities. So how do these models align with actual live test results?

As it turns out, they correlate pretty well with empirical data from our VHD Testbed. [Figure EC2-8](#) shows the per-client totals for the same client scaling data presented in [Figure EC2-2 on page 16](#). Whereas that figure shows the aggregate goodput from 10 to 100 clients, this chart shows the per-client goodput for the same test runs. Note that per-device throughput converges to approximately 300 Kbps for every spatial stream configuration except the pure 3SS client population case. As you have learned, it is not expected that 3SS clients will form a significant part of most VHD environments in the future.




Figure EC2-8 Per-Device TCP Throughputs from Aruba VHD Lab

This data was collected in a clean air environment with no interference. So if you apply an impairment factor, you get numbers that are very consistent with the system-wide estimate above.

Conclusions

In this chapter, you have learned how to estimate channel throughput in various conditions, to roll up multiple channels into a total system throughput (TST), and how to divide a channel across a given device population to estimate per-user performance.

It's clear that the worst case scenario for any VHD network is to compute a gross capacity of $X \text{ Mbps} * Y$ channels, and assume no spectrum reuse. This approach allows for the possibility of better than expected performance if spatial reuse is achieved while managing downside risk if it is not.

RF spatial reuse should always be assumed to be equal to 1 unless you can prove that it is higher. You must know that the design you plan to use is capable of achieving spatial reuse in that specific facility type. Spatial reuse is nearly impossible to achieve in facilities of 10,000 seats or less. If your deployment requires reuse to achieve your application requirements, consult your local Aruba systems engineer.

Chapter EC-3: Airtime Management

In the previous chapter, you learned how critical the average channel throughput parameter is to estimating the capacity available at the system and per-device levels. The throughput estimation process assumes that the WLAN configuration has been optimized to maximize idle airtime and minimize transmissions as much as possible. The Aruba VHD lab data presented in this VRD was obtained with an optimized configuration.

In this chapter, you learn how to configure an Aruba system for optimal airtime operation for very high-density (VHD) facilities. Aruba engineers have developed a best practice configuration template for these environments that has been proven in deployments worldwide.

Maximizing Capacity Through Good Configuration

ArubaOS can intelligently manage many aspects of VHD environment to provide the best possible experience to all users in the coverage area. To achieve this experience, the Aruba controller must be configured to continuously optimize the use of channels, clients, data rates, airtime, and low-level radio settings. Good configuration increases capacity by freeing up airtime. Poor configuration actually can destroy capacity by using airtime unnecessarily.

When learning VHD WLAN configuration, it is useful to think of these optimizations being applied in a specific sequence. The basic four-step approach is shown in the diagram. This chapter proceeds in that order.




Figure EC3-1 VHD Airtime Configuration Process

As always, this chapter assumes that a good RF design has been implemented that provides for solid 25-30 dB signal-to-interference-plus-noise ratio (SINR) throughout the VHD coverage area. If this level is not possible for some reason, then implementing the configuration suggestions covered here is even more

important to compensate for the lower data rates you will experience at lower SINR levels.

Good Client Distribution

To begin VHD configuration, ensure that clients are well distributed across the available channels. Ensure that virtually all 5 GHz-capable clients are moved to that band. Uneven client distribution is a major cause of poor VHD performance.

This section reviews these seven techniques:

- Send transmissions in parallel using only 20-MHz channels
- Use all 5-GHz channels
- Ensure a good channel plan
- Use power differentials to self-steer clients to the 5-GHz band
- Enable 802.11k and 802.11v for radio measurements and graceful steering
- Use QoS Basic Service Set (QBSS) load IE to self-steer clients to less busy APs
- Use ClientMatch for band steering and client load balancing

These approaches work together to help ensure that all the channels get as equal a load as possible. Equal loads reduce congestion and improve the overall user experience.

Send Transmissions in Parallel by Using Only 20-MHz Channels

For 802.11ac, Aruba is restating its longtime guidance that all VHD WLANs should only use 20-MHz channel widths. We have studied channel width in the lab, and it's clear that 20-MHz is the optimal channel width for large numbers of clients. Never use bonded channels in VHD areas, and here are four main reasons why.

- **Increased reuse distance** - Using 80-MHz (VHT80) or 40-MHz (VHT40) channels reduces the number of radio channels by bonding them together. Having fewer channels means that the distance between same-channel APs is also reduced because there are not as many channels to spread around. If you use 40-MHz channels, you have half the number of channels as 20-MHz, therefore each channel must be used twice as often. [Figure EC3-2](#) illustrates this concept.




Figure EC3-2 Adding Channels Increases Same-Channel Reuse Distance

Even if there was no performance difference between the channel widths, the reuse distance is a good reason not to use bonded channels. Larger reuse distances decrease ACI and CCI and improve overall system performance.

- **Reduced retries due to lower interference** – Wider channels double or quadruple the exposure of each and every transmission to interference from legacy devices transmitting with a smaller bandwidth. As a result, using narrow 20-MHz channels reduces retries, which allow some channels to get through even if others are blocked.
- **Higher SINRs** – 20-MHz channels experience up to 6 dB more SINR than 80-MHz channels for the same data rate, and up to 3 dB more SINR than 40-MHz channels. This increase is because the noise floor increases by about 3dB with each doubling of channel width. Higher transmit power is therefore required for any given MCS data rate, as compared with a narrower bandwidth. In practice, since many APs in VHD areas already operate at maximum power, the result is to reduce the SINR of the 40-MHz and 80-MHz signals.

This extra SINR can make a significant difference in the throughput of a VHD network. In essence, it provides additional margin in the link budget in both directions. This gain can help keep data rates higher when dealing with crowd loss, obstructions and other impairments.

- **Higher performance** – It is far better to have 25 users each on four different VHT20 channels than 100 users on one VHT80 channel. Four clients can be served at the same time in the same amount of spectrum, instead of having to be served consecutively by a single VHT80 access point.

The Aruba VHD lab evaluated this recommendation with 100 smartphones and 100 laptops. Each group of clients was tested with a single VHT80 AP, then split into two groups of 50 and tested with two VHT40 APs. Finally, they were split again into four groups of 25 clients and tested against four VHT20 APs. In all cases, we used the same 80 MHz of spectrum (channels 100 – 112).




Figure EC3-3 Test Design for VHT80 vs. VHT40 vs. VHT20

Figure EC3-4 shows the results for the 1SS smartphone for TCP upstream and bidirectional traffic. 20-MHz channels outperform 40-MHz channels, particularly above 50 stations. Four 20-MHz channels are nearly three times faster than one 80-MHz channel with 100 devices contending at the same time!




Figure EC3-4 Comparing VHT20 to VHT40 and VHT80 for 1SS Smartphone

Figure EC3-5 shows the results for a 2SS laptop. As seen with 1SS smartphone, the narrower 20-MHz channel pulls away beyond 25 concurrent transmitters. For 100 stations, the narrow 20-MHz channels produce almost 50% more throughput than 80-MHz for bidirectional traffic, and about 5% more for upstream traffic.




Figure EC3-5 Comparing VHT20 to VHT40 and VHT80 for 2SS Laptop

What is the explanation for this result? Shouldn't clients be served faster in a single 80-MHz channel because of the higher data rates?

When more stations contend for access to the channel, the efficiency of the MAC layer decreases due to an increase in protocol overhead. We call this phenomenon the “contention premium” and it is explored in depth in [Chapter T-4: How Wi-Fi Channels Work Under High Load](#) of the *Very High-Density 802.11ac Networks Theory Guide*.

As a general rule, for any large group of Wi-Fi® devices, it is better to divide them across more small channels to allow them to send in parallel.

The main benefit to using wider channels is the ability for individual stations to burst at the maximum PHY rate when only a portion of the users are trying to use the WLAN. However, in the VHD scenario, **we are not designing for peak burst rate**. Instead, we are designing to provide a consistent minimum bandwidth to all users (such as 512 Kbps or 1 Mbps). In this case, we accept a reduction in the maximum per-station burst rate during light loads in exchange for a greater total user capacity at all times.

To configure the Aruba system to only use 20-MHz channels, adjust these profiles as shown in [Table EC3-1](#). (These profiles are described in more detail in [Chapter EC-4: Channel and Power Plans](#).)

Table EC3-1 Profiles and Settings for Disabling Channel Bonding

Profile	Option	Setting	Reason
HT-SSID Profile	40-MHz-enabled	Disabled	Disables 40-MHz operation on the SSID.
	80-MHz-enabled	Disabled	Disables 80-MHz operation on the SSID.
ARM Profile	40-MHz-allowed-bands	None	Disables 40-MHz channel assignment. (Must be done on ARM profile for both radios.)
	80-MHz-support	None	Disables 80-MHz channel assignment. (Must be done on ARM profile for both radios.)
Regulatory Domain Profile	valid-11g-40mhz-channel-pair	No	Removes all 40-MHz channel pairs (enabled by default).
	valid-11a-40mhz-channel-pair	No	Removes all 40-MHz channel pairs (enabled by default).
	valid-11a-80mhz-channel-pair	No	Removes all 80-MHz channel pairs (enabled by default).

Use All 5-GHz Channels (Except 144)

As has been stated several times already, all available 5-GHz channels should be used in VHD areas, including DFS channels. It bears repeating here because the primary reason is to achieve good client distributions.

Using more channels directly spreads the client load out, which reduces the number of stations on each channel. Using more channels also reduces the MAC overhead, collisions, retries, and other impairments experienced when contention levels are high. As a result, both Total System Throughput and Per-Device Throughput are increased.



At the time of this writing, channel 144 should be used with care. Only the latest 802.11ac-capable devices can use it. No 802.11n or 802.11a clients will see this channel.

To configure the Aruba system to use only all channels, adjust this profile as follows. (This profile is covered again in [Chapter EC-4: Channel and Power Plans](#) in more detail.)

Table EC3-2 Profiles and Settings for Enabling All Channels

Profile	Option	Setting	Reason
Regulatory Domain Profile	valid-11a-20mhz-channel-pair	All allowed channels except 144	Ensures that all allowed channels in your country / regulatory domain are enabled. In some countries, DFS channels are disabled by default.

Ensure a Good Channel Plan

It does no good to activate all the available channels, if the channels are distributed poorly in the VHD area. Follow these guidelines to ensure a good channel plan:

- Do not repeat 5-GHz channels in a VHD area (unless reuse is explicitly part of the design).
- Evenly distribute 5-GHz channel numbers throughout the room (so adjacent APs are not on adjacent channel numbers).
- Ensure that 2.4-GHz channel numbers follow a regular repeating pattern.

Channel and power planning is a large and complex topic, and it is the main subject of [Chapter EC-4: Channel and Power Plans](#).

Use Power Differentials to Self-Steer Clients to 5-GHz

It is always better for a client to steer itself to the 5-GHz band on its own, than to have the WLAN infrastructure attempt to forcibly band-steer it. Infrastructure-led band-steering can be disruptive to clients, and clients can ultimately refuse. This type of band steering also generates extra low-rate control traffic that reduces overall capacity.

Aruba has found that many 5-GHz capable clients can be induced to prefer the 5-GHz band by using a 6 dB or 9 dB EIRP differential between the bands. So if your 5-GHz EIRP has been configured for 18 dBm, you would choose either 12 dBm or 9 dBm as the maximum EIRP in 2.4-GHz.

This technique is extraordinarily effective with clients that are available as of the time of writing, including legacy HT clients and newer VHT clients.

To configure power differentials, adjust the profiles shown in [Table EC3-3](#).

Table EC3-3 Profiles and Settings for Setting Up Power Differentials

Profile	Option	Setting	Reason
ARM Profile	max-TX-power	6 dB or 9 dB higher on the Dot11a ARM profile than on the Dot11g ARM profile	Enforces power differential. (Must be done on ARM profile for both radios.)
	min-TX-power	6 dB or 9 dB higher on the Dot11a ARM profile than on the Dot11g ARM profile	Enforces power differential. (Must be done on ARM profile for both radios.)

Use ClientMatch and 802.11v BSS Transition for Steering and Load Balancing

The capacity planning methodology we presented in [Chapter EC-2: Estimating System Throughput](#) assumes that the majority of 5-GHz capable clients are using that band. This methodology also assumes that within each band, the clients are evenly distributed across available APs and channels. ClientMatch is the ArubaOS feature that helps manage this automatically. ClientMatch was introduced in ArubaOS 6.3 and significantly enhanced in 6.4. ClientMatch replaces our legacy Band Steering and Spectrum Load Balancing features in earlier ArubaOS releases.

For those 5-GHz capable clients that still choose 2.4-GHz even after employing power differentials, use ClientMatch to attempt to force them to switch. This intervention is necessary because the drivers in many dual-band clients still have a preference for the 2.4-GHz band, even when 5-GHz service is available. To affect a switch, we use a combination of classic band steering techniques as well as new technologies based on 802.11v.

ClientMatch also provides a load-balancing service. The ClientMatch algorithm seeks to equalize the number of clients on each available channel. It has a variety of tools at its disposal to attempt to steer clients away from more heavily loaded APs. Though the decision of which AP to choose ultimately is left to the client, ClientMatch has been shown in lab tests to improve balance across radios in all types of deployments.

As part of 802.11v implementation, ArubaOS supports Basic Service Set (BSS) Transition Management. For clients that support 802.11v, ClientMatch uses BSS transition commands instead of 802.11 deauthentication commands to perform steering. This method is more graceful, less disruptive, and more likely to be successful than forcibly breaking the association of an existing device. 802.11v Fast BSS Transition requires that 802.11k be enabled.

ClientMatch is enabled by default in ArubaOS 6.3 and higher. No command is needed to turn it on. However, you should implement a set of ClientMatch optimizations for VHD environments.

Table EC3-4 Profiles and Settings for ClientMatch Optimization for VHD

Profile	Option	Setting	Reason
ARM Profile	cm-sticky-snr	18 dB	Minimum SNR to avoid being steered.
	cm-lb-client-thresh	50	Minimum client count on an AP before steering will occur.
	cm-lb-snr-thresh	20 dB	Min SNR of candidate AP in order to steer a client.
	cm-sticky-min-signal	-70 dBm	Minimum RSSI of candidate AP in order to steer a client.
	cm-band-g-max-signal	-10 dBm	2.4-GHz clients with strong RSSI should still be steered to 5 GHz.
	cm-steer-timeout	3 seconds	Number of seconds that non-candidate APs should ignore client being steered.
	cm-max-steer-fails	3	Maximum number of steer attempts before client is marked as "unsteerable".
	cm-unst-ageout	Enable	Enforce ageout of unsteerable client table.
	cm-unst-ageout-intvl	4 hours	Duration that unsteerable client state will be retained.
	cm-stale-age	600 seconds	How long APs hold on to VBR data.
	cm-dot11v	Enable	Helps clients to roam faster, without using 802.11 deauthentication frame. (Enabled by default in 6.4.2.3 and later.)

Enable 802.11k Radio Measurements

802.11k must be enabled in order to use 802.11v BSS transition management commands with ClientMatch.

The 802.11k protocol provides mechanisms for APs and clients to measure the available radio resources dynamically. In an 802.11k-enabled network, APs and clients can send neighbor reports, beacon reports, and link measurement reports to each other. These reports allow clients to take the appropriate action when there is an issue with the connection.

Along with 802.11k, ArubaOS supports Radio Resource Management Information Elements (RRM IEs).

Table EC3-5 Profiles and Settings for Enabling 802.11k

Profile	Option	Setting	Reason
Virtual AP Profile	dot11k-profile	<Profile Name>	Dot11k profile must be defined in the VAP profile.
Dot11k Profile	dot11k-enable	Enabled	Enables dot11k operation.
	bcn-measurement-mode	Active	Client should use active probing to populate BSS table.
	rrm-ie-profile	<Profile Name>	Specifies RRM subprofile.
RRM IE Profile	quiet-ie	No	Do not silence channel for measurement reports. Required for interoperability and because VHD channels should never be silenced.

Use QBSS Load to Self-Steer Clients to Less Busy APs

In a similar manner as band selection, it is better for clients to steer themselves to more lightly loaded APs. The 802.11k amendment allows APs to advertise their current traffic and available capacity via the QBSS IE in the BSS beacon. 802.11k-capable clients can incorporate this information into their AP selection algorithms.

Table EC3-6 Profiles and Settings to Enable QBSS IE

Profile	Option	Setting	Reason
SSID Profile	qbss-load-enable	Enabled	Enables the AP to advertise the QBSS load element, which includes: ● Station count: The total number of stations associated to the QBSS. ● Channel utilization: The percentage of time the channel is sensed to be busy. ● Available capacity: The remaining amount of medium time (measured as number of 32us/s) available for a station. A QBSS-enabled device uses these parameters to choose the best AP.

Maximize Data Rates

When the user devices are properly distributed across the available channels, the average MCS rate for frames with user data is maximized because the signal quality of each client is also maximized. However, we must also raise the data rates used for control and management frames and also multicast data traffic. Most frames on a WLAN are control frames that are sent at low rates by default, so if clients send them faster, more capacity is available. In this portion, we'll look at these seven techniques to maintain the data rates of various traffic types as high as possible:

- Enable VHT rates on the 2.4-GHz band
- Maximize the rate of 802.11a and 802.11n data frames by “trimming” low rates
- Increase the rate of control and management frames to at least 24 Mbps
- Disable 802.11b rates and 802.11b protection mode
- Use a high beacon data rate of at least 24 Mbps
- Use dynamic multicast optimization
- Use multicast rate optimization

Enable VHT Rates on 2.4-GHz

802.11ac is designed exclusively for use in the 5-GHz band, including the 256-QAM data rates. However, certain major radio manufacturers have adapted this functionality for the 2.4-GHz band. Aruba APs support this capability, as do smartphones and tablets from leading vendors. Therefore, it is a good idea to enable VHT rates in the 2.4-GHz band to speed up those clients that can handle it.

Table EC3-7 Profiles and Settings to Enable VHT in 2.4-GHz

Profile	Option	Setting	Reason
Dot11g Radio Profile	very-high-throughput-rates-enable	Enabled	Allows 256-QAM data rates in 2.4-GHz for compatible clients.

Maximize Rate of 802.11a and 802.11n Data Frames by “Trimming” Low Rates

If the RF design has successfully delivered 25-30dB SINR everywhere (after accounting for crowd loss or other impairments) then the lowest 802.11a data rates need not be permitted. In fact, these are three good reasons to eliminate the low rates:

- To shorten the transmit time required for 802.11 Null Data Packets (NDP) used for channel sounding and to signal power state transitions
- To force clients approaching a cell edge to roam sooner by eliminating choices from the client roaming algorithm
- To avoid rate-adapting down to slow rates in low-SINR conditions

Ensure that the lowest TX rate is one less than the beacon rate. The beacon rate should never be exactly the same as the lowest TX rate, to allow roaming clients to continue to communicate with the AP as they cross over the edge of the cell.

It is important to trim the low 802.11n HT rates in a similar manner as the legacy 802.11a rates. This change will ensure that HT stations cannot rate adapt down to the slowest rates. 802.11n allows an administrator to specify both a minimum and a maximum allowed MCS value. This change is made via the HT-SSID profile.



With 802.11ac, you cannot disable low MCS rates for data frames. In 802.11ac, we can only specify the maximum rate.

Table EC3-8 Profiles and Settings for Transmit Rate Trimming

Profile	Option	Setting	Reason
SSID Profile	a-tx-rates	18 24 36 48 54	Eliminates the 6 Mbps and 12 Mbps rates. Consider eliminating the 18 Mbps rate as well if your RF design will support it. Your lowest TX rate should be one lower than your beacon rate.
	g-tx-rates	18 24 36 48 54	Same as a-tx-rates.
HT-SSID Profile	supported-mcs-set	3-7,11-15,19-23	Eliminates MCS0 – MCS2 from the HT rate set

Increase Data Rate of Control and Management Frames

You have seen that control and management frames consume much more airtime than data frames. Aruba recommends using a minimum rate of 24 Mbps. Increase the control rate from the default of 6 Mbps to 24 Mbps to reduce the airtime consumed by these frames by 75%.

Unless you explicitly change it as recommended here, ArubaOS uses the lowest configured basic rate for most control and management traffic (except beacons). For some traffic types, ArubaOS will use a higher rate if configured and there is adequate SINR. However, do not use a control rate higher than 36 Mbps as this can make it more difficult for these frames to get through in a VHD area. You should test this configuration with your specific mix of client devices to ensure that users report no issues from drivers that expect a fuller set of rates. While not very common, it can happen.

Table EC3-9 Profiles and Settings for Basic Rate Trimming

Profile	Option	Setting	Reason
SSID Profile	a-basic-rates	24 36	Sets minimum control and management frame rate on 5-GHz band to 24 Mbps. Allow for up-rating to 36 Mbps where possible.
	g-basic-rates	24 36	Sets minimum control and management frame rate on 2.4-GHz band to 24 Mbps. Allow for up-rating to 36 Mbps where possible. Eliminates 802.11b rates (implied by not specifying any).

Disable 802.11b Rates and Protection Mode

To help improve performance on the 2.4-GHz side, Aruba recommends that you completely disable the 802.11b rates unless they are specifically required by the clients in your particular population.

Remove all 802.11b rates from the TX rate and the Basic Rate sets in the SSID profile. Also, you should disable 802.11b protection mode if you do not expect to support such clients. Even with the 802.11b rates removed, the AP can still be affected negatively if it detects an 802.11b client in the area. This situation is very common. Disabling 11b protection may further improve performance in 2.4-GHz. The necessary setting is shown in [Table EC3-10](#).

The configuration for removing the rates is already implied in the SSID profile changes in [Table EC3-9 on page 46](#) (if you do not include 802.11b rates in the list, they are eliminated).

Table EC3-10 Profiles and Settings for Disabling 802.11b Protection Mode

Profile	Option	Setting	Reason
Dot11g Radio Profile	dot11b-protection	Disabled	Disables protection for 802.11b clients.

Use a High Beacon Rate

ArubaOS allows you to choose a specific data rate to be used for beacons. Aruba recommends that this rate be at least 24 Mbps in VHD areas. In some cases, 36 Mbps may be desirable. You must have average SINRs of 25-30 dB after accounting for crowd loss in order to use these rates.

As with control rates, you must test this rate with your specific mix of client devices to ensure that users report no issues.



Increasing beacon rate is usually much more effective to reduce beacon load than lengthening the beacon interval. Intervals larger than 200 milliseconds can cause compatibility problems with some clients.

For perspective, let us calculate the airtime consumed by beacons given the data rate used and the number of APs on the same channel. The charts in [Table EC3-11](#) show the percentage of airtime consumed by four different beacon rates, assuming a 290-byte beacon payload. (These rates were calculated with the Beacon Calculator from Revolution Wi-Fi.)¹

Table EC3-11 Airtime Consumption by Various Beacon Data Rates

APs Per Channel	Number of SSIDs		
	1	2	3
6-Mbps Beacon Rate			
1	0.44%	0.89%	1.33%
5	2.22%	4.44%	6.67%
10	4.44%	8.89%	13.33%
15	6.67%	13.33%	20.00%
20	8.89%	17.77%	26.66%
25	11.11%	22.22%	33.33%
30	13.33%	26.66%	39.99%
35	15.55%	31.10%	46.66%
40	17.77%	35.55%	53.32%
45	20.00%	39.99%	59.99%
50	22.22%	44.43%	66.65%
18-Mbps Beacon Rate			
1	0.19%	0.38%	0.57%
5	0.95%	1.90%	2.86%
10	1.90%	3.81%	5.71%
15	2.86%	5.71%	8.57%
20	3.81%	7.62%	11.43%
25	4.76%	9.52%	14.28%
30	5.71%	11.43%	17.14%
35	6.67%	13.33%	20.00%
40	7.62%	15.23%	22.85%
45	8.57%	17.14%	25.71%
50	9.52%	19.04%	28.56%
24-Mbps Beacon Rate			
1	0.16%	0.32%	0.48%
5	0.80%	1.59%	2.39%
10	1.59%	3.18%	4.78%
15	2.39%	4.78%	7.16%
20	3.18%	6.37%	9.55%
25	3.98%	7.96%	11.94%
30	4.78%	9.55%	14.33%
35	5.57%	11.14%	16.71%
40	6.37%	12.73%	19.10%
45	7.16%	14.33%	21.49%
50	7.96%	15.92%	23.88%
36-Mbps Beacon Rate			
1	0.13%	0.26%	0.38%
5	0.64%	1.28%	1.92%
10	1.28%	2.56%	3.84%
15	1.92%	3.84%	5.76%
20	2.56%	5.12%	7.68%
25	3.20%	6.40%	9.59%
30	3.84%	7.68%	11.51%
35	4.48%	8.96%	13.43%
40	5.12%	10.23%	15.35%
45	5.76%	11.51%	17.27%
50	6.40%	12.79%	19.19%

1. Revolution Wi-Fi, <http://www.revolutionwifi.net/capacity-planner>. Reprinted with permission.

Consider a football stadium with approximately 1,000 APs audible in the main bowl and two Extended Service Set Identifiers (ESSIDs). Approximately 20 channels will be used in the main bowl on 5-GHz. Therefore, there will be an average of 50 APs per channel. In this example, two SSIDs and a default beacon rate of 6 Mbps yield an expected airtime load of 44.4%. This rate is **per channel**. Increasing the beacon rate to 24 Mbps reduces this load to just under 16%.

Table EC3-12 Profiles and Settings for Fixed Beacon Rate

Profile	Option	Setting	Reason
SSID Profile	a-beacon-rate	24 or 36	Sets beacon data rate on 5-GHz band.
	g-beacon-rate	24 or 36	Sets beacon data rate on 2.4-GHz band.

Enable Dynamic Multicast Optimization (DMO)

In 802.11, multicast frames must be sent at the lowest basic rate in the BSS. Using the system defaults, a 2SS VHT20 client that is capable of 173.3 Mbps rate would have to slow down to 6 Mbps to receive multicast video (a reduction of almost 30X).

The Aruba DMO feature speeds up this rate dramatically by converting multicast data into unicast traffic in real time, which can be transmitted at full MCS rates supported by each device. DMO applies to all downstream multicast traffic, including video and non-video alike.

Unicast traffic is acknowledged at the Wi-Fi MAC layer, so DMO also improves reliability because multicast traffic is not acknowledged. Lost unicast frames are detected and resent automatically by the Wi-Fi driver.

A breakeven point exists for the number of clients in a BSS above which multicast becomes more efficient than unicast. The limit varies by Wi-Fi generation, and faster 802.11ac stations are able to support more unicast streams than older 802.11a or 802.11n stations. It is not very common to hit this limit unless you are streaming video to many clients, which happens only in special use cases, such as lecture halls.

The breakeven point defaults can be configured by the administrator using the *dynamic-mcast-optimization-threshold*. Transmission automatically switches back to multicast when the client count increases above the breakeven point so that the efficiency of unicast is lost.



IGMP proxy should always be used with DMO to ensure that the wired infrastructure sends video traffic only to those APs that have subscribers. Aruba has deprecated IGMP snooping as of ArubaOS 6.X. IGMP proxy is strongly recommended.

Table EC3-13 Profiles and Settings for Dynamic Multicast Optimization

Profile/Interface	Option	Setting	Reason
Virtual AP Profile	dynamic-mcast-optimization	Enabled	Converts multicast traffic to unicast below the optimization threshold.
	dynamic-mcast-optimization-threshold	80	Specifies multicast subscriber count cutoff above which transmission reverts to multicast.
Interface VLAN <X>	ip igmp proxy	gigabitethernet <slot>/<port>	IGMP proxy is applied to the VLAN interface, which in turn references physical ports.

Enable Multicast-Rate-Optimization

For nonvideo multicast traffic that cannot be converted to unicast with DMO, we must increase the data rate, if possible.

The default behavior in 802.11 is to transmit multicast traffic at the lowest configured basic rate for the BSS, so it stands the best chance of reaching all associated clients. Further, multicast transmissions are not acknowledged in 802.11, thus multicast delivery is inherently unreliable. Retries use the same low data rate. This behavior can be very expensive in terms of time on the medium, and multicast has been the subject of many optimization techniques.

The Aruba MRO feature keeps track of the transmit rates that are sustainable for each associated client and uses the highest possible common rate for multicast transmissions. This feature can reduce the airtime required to send a multicast stream. The basic operation is shown in [Figure EC3-6](#).




Figure EC3-6 Optimization Methods for Multicast Rate Selection

The figure shows five different STAs in a BSS. Stations V, W, and Y are members of a multicast group. The lowest data rate in use by these three stations is 24 Mbps. Therefore, with MRO enabled, the Aruba system sends multicast traffic at 24 Mbps instead of a lower rate.

Table EC3-14 Profiles and Settings for Multicast Rate Optimization

Profile	Option	Setting	Reason
SSID Profile	mcast-rate-opt	Enabled	Activates dynamic multicast rate selection based on the highest "common" data rate.

Eliminate Unnecessary Transmissions

Airtime is such an extremely limited resource in VHD areas that only truly necessary frames should ever go on the air. Unnecessary frames not only consume airtime once, but some percentage fails and must be retried.

If you watch wireless traffic with a protocol analyzer in any high-density area, you will be amazed at the sheer volume of unnecessary transmissions. Every single transmission represents lost capacity. If you implement the tools discussed in this section, you will recover much of that capacity and improve the user experience:

- Eliminate SSIDs and beacons
- Drop multicast and broadcast to minimize “chatty” protocols
- Do not use IPv6
- Use AirGroup if multicast Domain Name System (mDNS) or Digital Living Network Alliance (DLNA) services are required
- Enable IGMP proxy
- Deny interuser bridging
- Set local probe response threshold
- Increase A-MSDU to 3 and enable end-to-end jumbo frames

Eliminate SSIDs and Beacons

Beacons are one of the most important frame types to eliminate in VHD areas. When additional ESSIDs are created for specific purposes such as faculty or ticketing or point-of-sale terminals, these must send beacons to be heard, which leaves less airtime for actual user traffic.

Beacons are sent 10 times **per second per ESSID per AP per radio**. Even with a beacon payload optimized to 24 Mbps, [Table EC3-11 on page 48](#) shows how the load increases linearly with each extra ESSID.

To maximize performance, it is essential to use the absolute minimum number of ESSIDs. Aruba strongly recommends using no more than two SSIDs in most VHD areas:

- An open ESSID for guests, usually with a captive portal
- A secure ESSID for staff and other permanent “house” devices

Aruba further recommends that the secure ESSID only be available on 5 GHz, which reduces the beacon load on the 2.4-GHz band. Limiting secure users to 5-GHz will also improve reliability of staff devices.

SSID design is a complex topic, and is the subject of [Chapter EC-5: SSIDs, Authentication, and Security](#).

Drop Broadcast and Multicast to Block “Chatty” Protocols

A “chatty” protocol is one that sends small frames at frequent intervals, usually as part of its control plane for service discovery. Such frames are most often sent using multicast or broadcast. These small frames are a poor use of scarce airtime, and they should be exterminated whenever possible unless they are part of actual data transmissions. Wherever chatty protocols are not needed, they should be blocked or firewalled. These protocols include IPv6 if it is not in production use, netbios-ns, netbios-dgm, Bonjour, mDNS, UPnP, and SSDP (among many others). You should enable the “Broadcast Filter All” feature in the Virtual AP profile for each ESSID unless it is absolutely necessary to support multicast or broadcast-based applications. This feature will automatically block chatty protocols.

The only exceptions are if multicast video or IPv6 is required.

In either of these cases, this feature cannot be enabled. If AirGroup is enabled, ArubaOS creates smart pinholes through the firewall to permit mDNS traffic for specific allowed services to specific hosts.

Table EC3-15 Profile and Settings for Broadcast Filter All

Profile	Option	Setting	Reason
Virtual AP Profile	broadcast-filter	All	Prevents all broadcast and multicast traffic from being transmitted on the VAP. (<i>AirGroup traffic is excepted on a station-by-station basis.</i>)

Do Not Use IPv6

IPv6 should not be enabled in VHD areas at this time unless absolutely necessary.

ArubaOS fully supports IPv6 operation, and if your environment requires IPv6, you should feel confident that it will work. However, running a dual-stack environment can generate significant IPv6 control-plane traffic. User applications can generate additional network traffic. Therefore, enabling IPv6 will likely reduce the overall airtime that is available to clients to some degree.

In addition, multicast is required for IPv6 operation, so you cannot use the broadcast-filter-all command recommended earlier. IPv4 multicast and broadcast traffic will no longer be filtered.

Stadiums, convention centers, airports, and other public venues generally have no need for IPv6 at this time. University lecture halls and certain enterprise or government customers that run significant IPv6 installations may have to weigh the pros and cons of offering IPv6 in VHD areas.

IPv6 is disabled by default on ArubaOS. As with IPv4, the Aruba controller should never be the default gateway in a mission-critical, clustered deployment.

Use AirGroup If mDNS or DLNA Services Are Required

In general, you should seek to avoid offering access to multicast-based media servers, print servers, and other zero-configuration network resources in VHD areas. The multicast reasons for this have already been explained.

Aruba recognizes that certain VHD environments like lecture halls may rely on Apple TVs, ChromeCast, or other media players, and printers that utilize mDNS or DLNA for discovery.

AirGroup is an mDNS proxy service. It leverages the firewall functionality of the Aruba controller to punch smart pinholes through the *broadcast-filter-all* block to permit individual devices to access these services in spite of a general block on broadcast and multicast traffic.

Enable IGMP Proxy

This feature was already described in the context of multicast video, but it is repeated here because IGMP proxy is a good idea to run more generally for any other type of multicast traffic.

When IGMP proxy is enabled, ArubaOS ensures that the wired infrastructure sends traffic to only those APs that have multicast subscribers. When this feature is not enabled, multicast traffic on each virtual AP VLAN is replicated on every radio.

Table EC3-16 Settings for Enabling IGMP Proxy

Profile	Option	Setting	Reason
n/a	interface VLAN <X> ip igmp proxy	gigabitethernet <slot>/<port>	IGMP proxy is applied to the VLAN interface, which in turn references physical ports.

Deny Inter-User Bridging

Aruba recommends against allowing peer-to-peer (P2P) applications in VHD environments for security reasons. In addition, airtime that is consumed by discovery protocols, as well as the P2P applications themselves, can significantly impact channel capacity for other users.

Table EC3-17 Profile and Settings for Inter-User Bridging

Profile	Option	Setting	Reason
Virtual AP Profile	deny-inter-user-traffic	Enabled	Prevents peer-to-peer communication.

Set a Local Probe Response Threshold

802.11 probe requests and responses are normal features of WLANs that can quickly spiral out of control in VHD areas with many APs. In addition to trimming low data rates, you should also configure a local probe response threshold (LPRT) threshold. This threshold is expressed in SINR, and it represents the minimum SINR at which an AP must receive a probe request to provide a probe response.

Aruba recommends beginning with a value of 6 dB. Never go higher than a maximum value of 10 dB. Higher values can cause problems for roaming devices. The LPRT setting must be 3 dB less than the ClientMatch *cm-sticky-snr* parameter.



As of ArubaOS 6.3, LPRT also governs the SINR required for an AP to respond to an 802.11 authentication request. Auth requests below the threshold are ignored, which can create problems for roaming clients if the LPRT threshold is set too high. Do **not** use a threshold higher than 10 without discussing first with your Aruba systems engineer.

Table EC3-18 Profile and Settings for LPRT

Profile	Option	Setting	Reason
SSID Profile	local-probe-req-thresh	6 to 10	Reduces probe response traffic.

Increase A-MSDU to 3 and Enable End-to-End Jumbo Frames

802.11ac increases the maximum allowed size of aggregated MSDUs (A-MSDUs) and improves standardization of this feature across clients. By using a larger A-MSDU value, the AP can pack more data into a single transmission, which reduces the number of TXOPs needed to send a given amount of data.

By default, the Aruba controller uses a value of 2. Increase to 3 for VHD environments.



Jumbo frames must be enabled on the wired infrastructure between the controller and the APs to use larger A-MSDU values. Increasing the A-MSDU size without enabling jumbo frames will increase fragmentation in the AP tunnels, which reduces performance. Jumbo frames are enabled by default on ArubaOS, and they are discovered automatically if the wired path allows them.

Table EC3-19 Profile and Settings to Increase A-MSDU Value

Profile	Option	Setting	Reason
HT-SSID Profile	max-tx-a-masdu-count-be	3	Allows 3 MSDUs per A-MSDU in [BE] queue.
	max-tx-a-masdu-count-bk	3	Allows 3 MSDUs per A-MSDU in [BK] queue.
	max-tx-a-masdu-count-vi	3	Allows 3 MSDUs per A-MSDU in [VI] queue.

Maximize Channel Accessibility

In this chapter you have learned how to achieve the optimal client distribution across APs, to maximize on-air data rates, and to eliminate unproductive transmissions. The fourth and final area of airtime optimization through configuration is to ensure that all stations have maximum access to the medium. This process includes these steps:

- Use QoS and configure differentiated service code point (DSCP) markings
- Use Airtime Fairness to prevent starvation of some clients by others
- Do not use bandwidth contracts
- Use Cell Size Reduction feature to filter low-SINR CCI

Use Quality of Service

If voice or video clients are expected in the VHD WLAN, quality of service (QoS) must be implemented in the air and on the wire, end-to-end between the APs and the media distribution infrastructure. Wi-Fi Multimedia (WMM) is enabled by default in 802.11n and 802.11ac, and all channel access is governed by access categories. Unmarked transmissions default to AC[BE].

If you are supporting voice and/or video in the VHD area, configure the DSCP-to-WMM mappings for the values in use. Verify that the mobile device applications are properly marking QoS frames being sent upstream to the WLAN. To verify, inspect the MAC header with a packet capture utility.

Table EC3-20 Profile and Settings for QoS

Profile	Option	Setting	Reason
SSID Profile	wmm	Enabled	Enables Wi-Fi Multimedia. (This is enabled by default in 802.11ac.)
	wmm-vo-dscp	56	Sets explicit DSCP-to-WMM queue mappings. Adjust as necessary if your network is using non-standard values. (These are not set by default.)
	wmm-vi-dscp	40	
	wmm-be-dscp	24	
	wmm-bk-dscp	8	

Use Airtime Fairness

Airtime fairness is a basic requirement of any VHD environment with an unpredictable and heterogeneous client mix. Older 802.11a/b/g/n clients that require more airtime to transmit frames must not be allowed to starve newer VHT clients. Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) minimizes collisions, but it does not require any type of fairness for clients that are associated at different data rates.

The ArubaOS Airtime Fairness (ATF) algorithm uses infrastructure control to manage airtime allocation dynamically for every client in a BSS. This algorithm considers the traffic type, client activity, traffic volume, and other factors before it allocates airtime on a per-client basis for all its downstream transmissions. This

allocation ensures that with multiple clients associated to the same radio, no client is starved of airtime and all clients have acceptable performance.

The time allocation policy has three options:

- **Default access:** Disables airtime allocation.
- **Fair access:** Allocates same air time to all clients by the process of token allocation (regardless of radio capabilities).
- **Preferred access:** Allocates more air time to VHT or HT clients.

Preferred access is generally recommended for VHD WLANs. This option applies higher weights to faster modes, for example, which assures that an 802.11ac client can complete a transaction much faster than its 802.11a equivalent. Preferential fairness offers the highest overall data capacity, but at some cost to less-capable clients. Some network managers welcome this as a subtle nudge to the user population to upgrade to 802.11ac clients.

A good way to evaluate the effect of the ATF feature is to examine its impact on individual clients and flows. The throughput vs. time graph from a 25-client TCP downstream Ixia Chariot test in [Figure EC3-7](#) shows the difference in the individual client throughput when the shaping policy is toggled off and on. On the left, ATF is disabled and on the right it is enabled. The quasi-random, contention-based access on the left gives way to a much steadier result on the right due to the airtime shaping algorithm.




Figure EC3-7 Effect of Airtime Fairness Algorithm on Individual Flows

To configure Airtime Fairness, adjust the profile in [Table EC3-21](#).

Table EC3-21 Profile and Setting for Enabling Airtime Fairness

Profile	Option	Setting	Reason
Traffic Management Profile	shaping-policy	preferred-access	ATF rate limits the air and is vital to prevent older, slower clients from starving newer, faster clients.

Do Not Use Bandwidth Contracts

It is not necessary to configure any type of bandwidth contracts (BWC) in a VHD area. BWCs are not needed for two reasons.

First, the Aruba ATF feature is effectively a form of BWC. By imposing fairness on the amount of time that each station gets the channel, it also implicitly caps the amount of bandwidth that is available.

Second, the air in VHD areas is self-throttling. You have learned how RF spatial reuse is almost never seen in VHD environments. VHD environments are often single collision domains. Therefore, the CSMA/CA process itself intrinsically caps how much throughput it is possible to achieve. And also as you have seen, the total system throughput (TST) drops as more stations are added.

Therefore, adding BWCs is unnecessary and simply burns up controller or AP CPU cycles for no incremental value.

Use the Cell Size Reduction Feature to Filter Low-SINR CCI

To reduce the effects of CCI, enable the Aruba Cell Size Reduction (CSR) feature with a low value.

CSR is a method of artificially adjusting the receive sensitivity of an AP radio, which helps the APs to reject interference from co-channel sources outside the high-density coverage area.

VHD areas have a large number of clients and APs at comparatively high signal strengths. Outside the VHD area, more APs and clients are at comparatively low signal strengths. However, because of how far these signals can travel and still be decoded, they can reduce the capacity of the VHD zone. Filtering such signals may result in the AP seeing the channel as idle more often.

The receive sensitivity of a Wi-Fi device, in this case an AP, defines the lowest signal level at which it can successfully decode a frame on the air. Receive sensitivity tuning can be used to fine tune the APs to “ignore” frames below a desired signal level. This tuning helps to reduce network degradation to outside interference.

CSR works by setting a threshold measured in dB. This threshold is a delta from the current noise floor. For example, a CSR threshold of 10 dB would cause the AP to ignore any transmission of -80 dBm or less if the noise floor is currently -90 dBm. Changing noise floors affects the sensitivity threshold. This effect can be visualized in [Figure EC3-8](#).




Figure EC3-8 Operation of Cell Size Reduction Threshold

When the CSR threshold is set higher, signals from more distant APs are ignored, transmit deferrals are reduced, and network throughput is increased.



Cell Size Reduction should not be confused with Receive Sensitivity Channel Reuse (RSCR). CSR was introduced in ArubaOS 6.3. RSCR is still available in ArubaOS but it should never be used.

Figure EC3-9 shows how CSR works. The data was collected from an upstream UDP test. The X-axis values are path loss in dB between the client and the AP. When CSR is set to 5, the AP can still receive frames with attenuation of ~113 dB from the client. The SINR of the frames received at this attenuation will be ~ 5dB and therefore the AP can still receive 6 Mbps frames. When CSR is set to 10, the receiver stops receiving at attenuation of ~102 dB (that is, the receive range of the AP will be reduced by ~10 dB). As the CSR value increases, you can see the AP-client communication is cut off below that threshold.




Figure EC3-9 Path Loss Test Result Demonstrating CSR Operation

Aruba recommends beginning with a value of 6 dB. The CSR threshold must **never** be set higher than about 10 dB. The CSR radius must never be farther than the roaming threshold radius of most client devices. If this happens, the AP stops responding to clients that think they are part of a valid BSS before they have a chance to roam away gracefully. This behavior can produce all kinds of destructive effects on the channel.¹ Figure EC3-10 shows this situation, with the client able to hear the AP, but the AP is unable to hear the client.




Figure EC3-10 CSR Threshold Must Never Be Set Higher Than Roaming Threshold

The figure uses absolute RSSI expressed in dBm on the vertical axis because client roaming thresholds are generally expressed in RSSI. In the example, a CSR threshold of 30 dB results in a receive sensitivity cutoff at -60 dBm given a noise floor of -90 dBm. Because this is well above the client roaming threshold of -75 dBm, the high CSR setting can cause unexpected client problems and significant unnecessary transmissions.

Table EC3-22 Profile and Setting for Enabling Airtime Fairness

Profile	Option	Setting	Reason
Dot11g Radio Profile	cell-size-reduction	6 – 10	Increase CCA idle frequency by filtering low-SINR PLCP preambles and data payloads.
Dot11a Radio Profile			

1. Chuck Lukaszewski and Liang Li, "14/1416r1 Observed protocol violations caused by Dynamic Sensitivity Control with roaming STAs," IEEE 802.11 TGax, September 2014

Chapter EC-4: Channel and Power Plans

In the previous chapter, you learned the importance of good client distribution across the channels used by the very high-density system. If some channels are more heavily loaded than others, users on those channels have less airtime and consequently a worse user experience.

Good client distribution requires a good channel plan. A channel plan is the mapping of channel numbers on to specific radios. It is assumed the reader is familiar with this concept from other types of WLAN deployments.

This chapter describes six major topics:

- Unique requirements for channel and power planning in VHD requirements
- Understanding the different types of channel plans
- How Aruba Adaptive Radio Management™ (ARM™) makes channel and power assignment decisions
- How to design channel plans for various types of VHD areas
- How to design power plans for VHD areas
- How to configure dynamic and static plans using Aruba ARM

How VHD Channel Requirements Differ from Conventional WLANs

Conventional WLAN deployments use a cellular structure, with a single access point (AP) to serve each cell. Cells are honeycombed or otherwise staggered to ensure even spacing, as shown in [Figure EC2-6 on page 28](#). APs are at least 15 m (50 ft) apart if not more, and they are often separated by walls, floors, or other structural materials. This separation provides for significant path loss between all APs in the system (at least 75 dB in open areas and over 90 dB in walled areas). The resulting high level of signal differentiation allows automated radio management systems like Aruba ARM to create dynamic channel plans.

VHD deployments are different in some key respects that impact the approach used for channel planning and the effectiveness of automated radio algorithms.

- Multiple, stacked radios serving the same physical area / cell
- Very small AP-to-AP separation, as little as 2-5 m (6-16 ft)
- Clear line of sight (LOS) between APs

As a result, most of the APs in the VHD area hear one another at strong, very similar power levels. AP-to-AP signal-to-interference-plus-noise ratio (SINR) levels are at least 10 dB and sometimes over 40 dB hotter than conventional layouts. This fact creates challenges for automated algorithms that depend on highly differentiated SINR values between APs. As the number of APs in the same RF collision domain increases, the problem gets progressively more difficult.

Facilities with Multiple Adjacent VHD Zones

Lecture halls, convention centers, and movie theaters may have multiple adjacent VHD areas that need to reuse some of the same channel numbers. However, no single VHD room should use the same channel number more than once if at all possible. It may be necessary to reserve a “house” or presenter channel to ensure smooth network access for speakers.

Stadiums and Arenas

The largest VHD environments are stadiums and arenas. These locations have other interesting differences from conventional WLAN deployments:

- Multiple, nested VHD service areas (concourses, suites, bowl)
- Discrete channel sets are used in different parts of the facility
- Some areas may be outdoors, which limits the channels that can be used
- Possibility of reserved channels for ticketing, POS, or IT systems
- Users are present only for limited amounts of time
- Significant turnover in users from event to event
- Continuous, massive levels of Bluetooth and “my-fi” hotspot interference on 2.4-GHz band

All of these factors affect how the wireless architect approaches the selection of channel plans.

Responding to Interference

One subtle but very important difference is how the WLAN infrastructure should respond to interference.

In a conventional deployment, when a new interference source is detected that degrades channel quality, the system should steer around it by choosing a better channel. This action can cause ripple effects but because APs are widely separated, such changes have limited impact on the rest of the system. Such changes can be configured to occur only during idle periods, to ensure no service impact to users.

In a VHD deployment, things are very different:

- Most if not all available channels are in use in the same physical space, so no free alternate channel is available.
- Interference is almost always higher during events; how do you decide what constitutes a harmful level?
- Changing a channel during an event disrupts service to users.
- Changing a channel during an event shunts users to remaining channels, which reduces capacity for them.
- Interference is often transitory, and only for the duration of the event.

Finally, the WLAN architect has carefully planned the entire channel plan. The architect has knowledge that the WLAN infrastructure does not.

So the WLAN architect has to weigh the costs and benefits of using the infrastructure channel management features in a way that is not at all necessary in conventional WLAN deployments. Steering around interference is often not desirable for VHD deployments. Often it is more important to maintain the integrity of the channel plan and allow individual clients to make their own decisions if their RF links deteriorate.

Types of Channel Plans

We have established that channel assignment for VHD areas must be approached in a different way than traditional WLANs. Now consider the different types of channel plans.

All channel plans can be categorized according to three criteria:

Dynamic vs. Static	Global vs. Local	Repeating vs. Non-repeating
• A dynamic channel plan is one that can change in response to external events, such as interference or system load. • In a static channel plan, the channel numbers are fixed and should not change.	• A global channel plan uses the same channel list for all APs that terminate on the system. • A local channel plan uses different channel lists for different groups of APs on the same system.	• A channel plan is repeating if the same channel number is used more than once in the same coverage area. • A channel plan is non-repeating if it cannot reuse the same channel number in the same coverage area.

Most indoor WLANs use the simple case of automated radio management with every allowed channel. In these terms, the channel plan of such WLANs is dynamic, global, and repeating.

But in VHD environments, you can leverage the three categories to create a rich variety of channel plans to suit different requirements. Some typical examples are these plans:

- Global and Non-Repeating
 - University building with multiple adjacent lecture halls, with all channels available but no individual lecture hall can use the same channel more than once. Could be static or dynamic depending on exact configuration method used.
- Local and Repeating
 - Outdoor stadium with the bowl area using outdoor-only channels and the suites and concourses using indoor-only channels.
 - Large arena bowl area with fixed channel assignments that uses DFS channels for capacity, while only non-DFS channels are in use in other parts of the facility.
- Static, Local, Non-Repeating
 - A concert hall with two dedicated ticketing APs at entry gates on channels 36 and 149. Every gate has an identical setup.
 - Stadium press box with four dedicated APs that are hard-coded to channels that cannot be used in the bowl seating area.
 - A convention center with a “house” channel dedicated to presenters that exists on one AP in the front of every individual meeting room (and nowhere else).
- Dynamic, Global, Non-Repeating
 - A press area with six APs that can use any channel but no channel can be used more than once

The only VHD environments in which you can use the usual indoor “hands off” approach are standalone facilities of a few thousand seats. For all other VHD environments, you will ensure higher system capacity and a better user experience by exercising fine-grain control over channel selection.

Choosing Dynamic vs. Static Channel Plans

VHD areas almost always require static channel plans.

The lecture hall or convention center case with multiple adjacent VHD service areas requires a non-repeating plan. However, because ARM does not have any knowledge of walls or building structure, it

cannot guarantee that channel numbers are not repeated in a single room. A static plan is the only way to achieve this guarantee.


In a large arena or stadium with many more APs than channels and heavy use of directional antennas, static plans are also mandatory.

In the Aruba system, you can configure ARM to leverage a static plan on the 2.4-GHz band to create a dynamic plan on the 5-GHz band. This configuration saves a lot of work and you learn how to do it in this chapter. We call these plans *automatic static channel plans*. These plans are extremely important.

You can create manual static channel plans using AP-specific configuration commands. This method works for smaller VHD facilities, but it is too much work in stadiums and so ARM is very helpful.

Requirements of Good VHD Channel Plan

It does no good to activate all the available channels, only to have channels unnecessarily repeated and/or poorly distributed in the VHD area. It is essential that your channel plan meet each of these requirements:



To make best use of scarce spectrum, we must optimize the distribution of RF spectrum to APs and clients. In any VHD WLAN, we want use as many allowed RF channels as possible, and ensure that they are properly distributed within the coverage area after accounting for in-band 802.11 and non-Wi-Fi transmissions outside the area.

Should You Disable “Surplus” 2.4-GHz Channels?

The 5-GHz band offers as many as eight times more channels than the 2.4-GHz band, so you must decide whether to disable some of the 2.4-GHz radios. To protect the very limited airtime in 2.4 GHz, beacons should not be sent from any more APs than absolutely necessary to absorb the associated device capacity (ADC) target for that frequency band.

This question applies only to VHD areas with many stacked APs that serve the same physical area, such as a lecture halls or ballrooms. In this case, you can use the dimensioning analysis to decide how many 2.4-GHz radios must be active. This number is based on the ADC and the expected breakdown of users between bands. You can disable the remainder by converting them into air monitors.

For larger deployments using directional antennas or picocells, often just one primary AP serves a given cell (even though many other APs can be heard on the same channel). In this case, we cannot disable any 2.4-GHz radios because each AP must provide dual-band service to each cell. A disabled 2.4-GHz radio would result in a coverage gap in those cells. The Aruba recommendation in [Chapter EC-2: Estimating System Throughput](#) to employ 6 – 9 dB power differentials between the bands further means that disabling 2.4-GHz radios is potentially harmful to users.

Aruba recommends that 2.4-GHz and 5-GHz cell sizes always be matched from an antenna pattern perspective.

How VHD Power Requirements Differ from Conventional WLANs

We have seen that conventional WLANs have one AP per physical cell, and the cells are honeycombed. The cell “edge” is commonly defined as the distance at which a minimum signal level is reached, such as -65 dBm. The AP-to-AP distance is different in every deployment, so an important responsibility of automatic radio management systems is to adjust the equivalent isotropic radiated power (EIRP) level to achieve the targeted minimum signal level in every cell.

By contrast, VHD deployments are complex, three-dimensional systems that stack multiple APs together to serve the same physical area. Even in stadiums and arenas where a single AP can be targeted at a specific block of seats, multiple APs on every channel hear one another at comparatively high SINR. Therefore, the goals and logic behind power selection in VHD areas are completely different from normal WLANs.

Another vital feature of VHD areas is large numbers of human bodies. RF signals are highly attenuated as a result. This effect is called “crowd body loss” or simply “crowd loss”. There are two mechanisms of action:

- **Reflection dampening:** With a large crowd present, multipath reflections are significantly reduced even in line-of-sight (LOS) conditions. Signals that might have bounced off a floor, wall, or empty seat are now absorbed or otherwise minimized. The intended Wi-Fi receivers are less able to recover extra spatial streams or to apply processing gain techniques like MRC. Aruba has observed typical reductions of 3 – 6 dB in overall SINR levels in certain VHD areas.
- **Direct absorption:** When the path between AP and client is directly blocked by human bodies, radio signals are significantly attenuated. This attenuation happens in all three RF coverage strategies explained in [Chapter P-3: RF Design](#) of the *Very High-Density 802.11ac Networks Planning Guide*. For example, in the overhead case it is common to shoot from behind through the body of a user that is facing away from the AP. Picocell configurations intentionally shoot through multiple bodies from underneath. Attenuation is cumulative with size of the crowd being traversed, and it can be modeled in a similar way to foliage loss outdoors.

In both cases, SINRs can be reduced dramatically as compared with an empty facility. Overhead loss due to crowd effect is typically on the order of 6 – 10 dB. Picocell crowd loss has been measured between 15 – 30 dB depending on the exact geometry of AP and client.

Facilities with Multiple Adjacent VHD Zones

Where multiple adjacent VHD rooms exist with the same channel number in each room, the idea of a “minimum cell edge” signal strength is quite irrelevant. It would be completely counterproductive for two relatively close same-channel APs separated by a wall to reduce EIRP, which could reduce the speeds for users in each room.

Stadiums and Arenas

An important area of difference is the use of external antennas in larger facilities like arenas, convention centers, and stadiums. These facilities increase the EIRP of the AP, which requires the wireless architect to decide separately how much conducted power should be directed to the antenna ports.

In addition, these facilities often use completely different combinations of APs and antennas to serve different area types. Concourse areas use a different antenna than the bowl, which is different again from skyboxes and club areas.

Requirements of Good VHD Power Plan

Power planning is therefore especially critical in all VHD areas. Just as with channel selection, it is virtually impossible to expect any automated radio algorithm to make the correct power level decisions in VHD areas. Due to the complex, 3D nature of VHD areas, a wireless architect simply cannot impart the intent and knowledge behind the design to the system.

Power Plans Should Deliver Client SINR Equal or Greater than 30 dB

The dominating concern of the wireless architect is to ensure that a minimum 25 – 30 dB SINR is available everywhere in the service area (after accounting for signal loss due to human bodies).



The importance of this criterion cannot be overstated. Forget everything you know about -65 dBm cell edges when you plan VHD areas.

The two reasons for this criterion are simple:

1. To maximize the modulation and coding scheme (MCS) for every client to free up airtime (capacity)
2. To stop client devices from probing and roaming unnecessarily

You have already read about the huge, negative effect on capacity by clients sending at slow PHY rates. WLAN architects should also read [Chapter T-3: Understanding Airtime](#) of the *Very High-Density 802.11ac Networks Theory Guide* for deeper insight in this area.

With respect to probing, most Wi-Fi clients reduce their management traffic significantly when they have a high SINR to their associated AP. Conversely, when SINR drops, most clients increase their probe activity in anticipation of roaming. Many clients are programmed to increase the aggressiveness of this behavior as the SINR decreases.

Management frames are sent at much lower data rates than data frames, and require full CSMA/CA arbitration which is expensive in terms of airtime. Therefore, hundreds or thousands of probing clients is a direct threat to the overall capacity of the system.

With 802.11ac, the MCS 9 data rate requires over 30 dB of SINR to be used. With 802.11n, the MCS 7 peak data rate requires approximately 25 dB of SINR.

Your cells must therefore target 25 – 30 dB of SINR **everywhere in the cell**. This SINR is net of human body loss. When you stack APs, every AP that cover the same area should be designed to deliver this level at the client.

Use the Absolute Minimum EIRP Possible in Each VHD Area

You must use the lowest EIRP possible to achieve the 25 - 30 dB SINR objective. Required minimum EIRP varies based on the type of antenna used by each AP, the distance between the antenna and the served users, structural losses, and other factors. You will have some APs where full power is necessary to achieve the minimum client SINR, and other APs where a lower power level will do the job.

This approach is important because:

- VHD areas have many multipath reflections (even after crowd effects), and using unnecessarily high power raises the overall noise floor for all users of that channel.
- If RF spatial reuse is being attempted, unnecessary power reduces the likelihood of success.
- Concourses, hallways, and skyboxes should minimize interference with meeting rooms and seating bowl areas.

The minimum EIRP level varies by RF coverage strategy and whether external antennas are in use. This decision requires that the wireless architect exercises proper judgment based on the particulars of the facility.

Power Should Be Measured by 1SS Devices

The minimum SINR requirement applies especially to single spatial stream devices such as smartphones, ticket scanners, point-of-sale terminals, and the emerging class of wearable computing devices.

Multistream devices such as phones, tablets, and laptops will see an AP at 1.5 – 3 dB higher than a 1SS device (at least). This slight increase in gain is due to maximal ratio combining (MRC) technology used by 802.11 on devices with multiple receive chains. As well as to improved antenna quality and positioning inside the devices.

The 1SS measurement requirement has important consequences for post-install validation. Many WLAN engineers use site survey utilities that run on PCs with multistream network adapters. These utilities are not acceptable because they do not provide a true picture of the environment. Several mobile site survey applications are now available that run on smartphones and tablets. Be sure to use a 1SS device with such software to validate your coverage. Also, be sure to test while a crowd is present.

Use a 6 – 9 dB Lower Power on the 2.4-GHz Band

As you learned in the last chapter, Aruba has found that most dual-band clients can be induced to prefer the 5-GHz band by using a 6dB or 9dB EIRP differential between the bands. So if your 5-GHz EIRP has been configured for 18 dBm, choose either 12 dBm or 9 dBm as the maximum EIRP in the 2.4-GHz band.

This technique is extraordinarily effective with clients that are available as of the time of writing, including legacy HT clients and newer VHT clients.

Use a Static EIRP Setting Based on Post-Install Tuning

Finally, use either a very narrow min/max range or a single static value for every radio in the VHD area. If you use a range, keep it to no more than 3 dB between the min and max value in the ARM profile.

After the system has been built, each different type of VHD area should be tested at various EIRPs, with and without a crowd present. From this data, you can “dial in” the correct power level to deliver the 25 – 30 dB SINR.

If you have multiple instances of the same type of room, you need to test only one or two power settings. If you have several different kinds of areas, such as concourses, skyboxes, and a seating bowl, evaluate each one.

This post-install tuning is mandatory for all VHD deployments and the associated labor must be factored into any service estimates.

How Aruba Adaptive Radio Management Makes Decisions

The Aruba ARM technology uses a distributed channel reuse management algorithm where each AP makes decisions independently by sensing its environment and optimizing its local situation. The algorithm is designed so that this iterative process converges quickly on the optimum channel plan for the entire network, but without a central coordinating function.

Background Scanning

Every 10 seconds, every Aruba AP leaves its home channel very briefly and performs a scan of another channel in its allowed channel list. These scans look for other APs, clients, rogue APs, background noise, and interference. Off-channel scans last for approximately 85 milliseconds, which provides enough time for the AP to get back to its home channel in time to send the next beacon. When the end of the channel list is reached, the AP starts over at the top.

Off-channel scanning is automatically deferred if the AP is busy. This behavior is configurable based on the level of load or the type of traffic.

Radio Metrics

Two metrics are maintained for every channel on every AP: the “coverage index” and the “interference index”. These indices are used to calculate the optimum channel as well as transmit power for the AP.




Figure EC4-1 ARM Channel and Transmit Power Selection Algorithm

The coverage index comprises the number of APs transmitting on a particular channel, weighted by their signal strengths as measured by the AP. The ARM algorithm aims to maximize and equalize coverage indices for all channels. This metric is the primary factor that controls AP transmit power within configured limits. ARM also seeks to maximize the separation of adjacent channels when possible, for instance separating channel 36 and 40 by at least one cell.

Interference index is a single figure that represents Wi-Fi activity and non-Wi-Fi noise and interference on a channel. When the interference index on the current channel is high compared to other channels, the AP looks for a better channel. Generally ARM chooses the channel with the lowest interference index. This choice avoids non-Wi-Fi interference, but also minimizes CCI as other APs on the same channel contribute to the interference index.

The Regulatory Domain Profile

The list of channels from which ARM can choose is determined by three factors:

- The *regulatory domain profile* in the AP group to which each AP is assigned
- The regulatory approvals Aruba holds for the specific AP model being used
- The indoor / outdoor setting assigned to the AP during provisioning

The *regulatory domain profile* is rarely if ever changed in Aruba deployments, so you may not be familiar with it. The default profile is configured on bootup from the country code set on the controller. DFS channels are disabled by default in certain countries including the US. For example, the default profile in the US regulatory domain is shown here.

```
(ArubaCLI) #show ap regulatory-domain-profile default
```

```
Regulatory Domain profile "default"
```

Parameter	Value
Country Code	US
Valid 802.11g channel	1
Valid 802.11g channel	6
Valid 802.11g channel	11
Valid 802.11a channel	36
Valid 802.11a channel	40
Valid 802.11a channel	44
Valid 802.11a channel	48
Valid 802.11a channel	149
Valid 802.11a channel	153
Valid 802.11a channel	157
Valid 802.11a channel	161
Valid 802.11a channel	165
Valid 802.11g 40MHz channel pair	1-5
Valid 802.11g 40MHz channel pair	7-11
Valid 802.11a 40MHz channel pair	36-40
Valid 802.11a 40MHz channel pair	44-48
Valid 802.11a 40MHz channel pair	149-153
Valid 802.11a 40MHz channel pair	157-161
Valid 802.11a 80MHz channel group	36-48
Valid 802.11a 80MHz channel group	149-161

As you will learn in the next section, the regulatory domain profile is a powerful tool to implement custom VHD channel plans.

Now, just because a channel is listed in the regulatory domain profile does not mean that it is available for use. If the AP model you are using does not have approval in your country, then ARM is not able to use

that channel. You see this when you enable DFS channels if the AP model has not yet received DFS approvals in a certain country.

Finally, channels that are marked as “indoor” only are not available to APs that are provisioned as “outdoor”. The opposite is also true. For example, channel 36 in the US cannot be selected on an AP that is set to outdoor.

ARM Profile

WLAN engineers should familiarize themselves fully with the *ARM profile*. This profile controls many different aspects of ARM decision making. The profile includes ClientMatch settings that you saw in [Chapter EC-2: Estimating System Throughput](#). Each radio on an AP has a separate ARM profile.

Configuring Dynamic and Static Channel Plans

A dynamic channel plan is one that can change in response to external events. To obtain a dynamic plan, set the ARM operating mode to *single-band*. This is the default mode of operation in ArubaOS.

In a static channel plan, the channel numbers are fixed and do not change. Set the ARM operating mode to *maintain* to implement a static plan.



This guide assumes that you are familiar with administration of ArubaOS via web browser and/or command line. The GUI and CLI configuration examples are based on ArubaOS 6.4.2.3. We are highlighting the important configuration steps, not providing step-by-step procedures.

Configuring the ARM Assignment Mode

The *assignment* parameter is used to set the operating mode. There are four *assignment* options – *single-band*, *multi-band*, *maintain* and *disable*. We are concerned with *single-band* and *maintain* for VHD deployments.



Separate ARM profiles must be configured for the 2.4-GHz and 5-GHz radios.

Table EC4-1 Profiles and Settings for Setting ARM Assignment Mode

Profile	Option	Setting	Reason
ARM Profile	assignment	single-band	Configures ARM for dynamic channel plan operation.
		maintain	Configures ARM for static channel plan operation.

Initial Channel Assignment for Static Plans

ARM cannot be placed directly into *maintain* mode. Allow for a training period during which ARM can create an initial channel plan.

All VHD deployments should put ARM into *single-band* mode for a brief period of time. The minimum training period for ARM to complete several scans is one hour. After that time, you can set *assignment* to *maintain* if you intend to run with a static plan.

During this period, temporarily disable three of the “-aware” knobs in ARM. At the end of the training period, enable these settings. This will accelerate the training period if the system is being used by test clients or real users.

Table EC4-2 Profiles and Settings for ARM Training

Profile	Option	Setting	Reason
ARM Profile	client-aware	Disable	Allows AP to change channel even if clients are associated to AP and passing traffic.
	video-aware-scan	Disable	Allows AP to go off-channel to scan even if active video clients are present.
	voip-aware-scan	Disable	Allows AP to go off-channel to scan even if active voice users are present.

Configuring Global and Local Channel Plans

A global channel plan uses the same channel list for all APs. With a global plan, all AP groups on the controller reference the same *regulatory domain profile*.

A local channel plan uses different channel lists for different AP groups. To implement a local plan, you must create unique *regulatory domain profiles* for each channel set. You then reference the appropriate one from each AP group.

Setting Up a Regulatory Domain Profile

Whether you will use a global or local plan, the first step is to create and modify a *regulatory domain profile*. All such new profiles are copied from the default profile. As a general rule, never modify the default profiles on the controller.

These profiles are most likely new to all readers, so we present GUI screen shots and CLI output to make the process more clear. For instance, [Figure EC4-2](#) shows the default *regulatory domain profile* for the US.

Configuration > AP Group > Edit "default"

Profiles

- Wireless LAN
- RF Management
- AP
 - Ethernet interface 0 port configuration: default
 - Ethernet interface 1 port configuration: default
 - Ethernet interface 2 port configuration: shutdown
 - Ethernet interface 3 port configuration: shutdown
 - Ethernet interface 4 port configuration: shutdown
 - AP system: default
 - Regulatory Domain: default**
 - Provisioning
 - AP authorization
- QoS
- IDS
- Mesh

Profile Details

Regulatory Domain profile > default Show Reference Save As Reset

Country Code	US - United States		
Valid 802.11g channel	☒ 1 ☐ 4 ☐ 7 ☐ 10	☐ 2 ☐ 5 ☐ 8 ☒ 11	☐ 3 ☒ 6 ☐ 9
Valid 802.11a channel	☒ 36 ☐ 60 ☐ 104 ☐ 116 ☐ 140 ☒ 153 ☒ 165	☒ 40 ☐ 64 ☐ 108 ☐ 132 ☐ 144 ☒ 157 ☒ 161	☒ 44 ☐ 56 ☐ 100 ☐ 112 ☐ 136 ☒ 149
Valid 802.11g 40MHz channel pair	☒ 1-5 ☐ 4-8 ☒ 7-11	☐ 2-6 ☐ 5-9	☐ 3-7 ☐ 6-10
Valid 802.11a 40MHz channel pair	☒ 36-40 ☐ 60-64 ☐ 132-136 ☒ 157-161	☒ 44-48 ☐ 100-104 ☐ 140-144	☐ 52-56 ☐ 108-112 ☒ 149-153
Valid 802.11a 80MHz channel group	☒ 36-48 ☐ 132-144	☐ 52-64 ☒ 149-161	☐ 100-112

Figure EC4-2 Default Regulatory Domain Profile in ArubaOS 6.4.2.3

The default profile must have these changes made to it for VHD deployments, which use exclusively 20-MHz channel widths:

- Add DFS channels if desired (except for channel 144 at this time).
- Remove 40 MHz channel pairs for 2.4-GHz and 5-GHz bands.
- Remove 80 MHz channel pairs.

In the GUI, the resulting view after making these changes should look like [Figure EC4-3](#).

Valid 802.11a channel	☒ 36	☒ 40	☒ 44
	☒ 48	☒ 52	☒ 56
	☒ 60	☒ 64	☒ 100
	☒ 104	☒ 108	☒ 112
	☒ 116	☒ 132	☒ 136
	☒ 140	☐ 144	☒ 149
	☒ 153	☒ 157	☒ 161
	☒ 165		
Valid 802.11g 40MHz channel pair	☐ 1-5	☐ 2-6	☐ 3-7
	☐ 4-8	☐ 5-9	☐ 6-10
	☐ 7-11		
Valid 802.11a 40MHz channel pair	☐ 36-40	☐ 44-48	☐ 52-56
	☐ 60-64	☐ 100-104	☐ 108-112
	☐ 132-136	☐ 140-144	☐ 149-153
	☐ 157-161		
Valid 802.11a 80MHz channel group	☐ 36-48	☐ 52-64	☐ 100-112
	☐ 132-144	☐ 149-161	

Figure EC4-3 Modified Regulatory Domain Profile with Only 20-MHz Channels and DFS

Here are the equivalent CLI commands:

Add DFS channels. (Depending on your country, these may already be enabled by default.)	<pre>! ap regulatory-domain-profile <profile name> valid-11a-channel 52 valid-11a-channel 56 valid-11a-channel 60 valid-11a-channel 64 valid-11a-channel 100 valid-11a-channel 104 ... valid-11a-channel 140 !</pre>
Remove 40-MHz and 80-MHz channel pairs. (The exact list of commands does vary from country to country; be sure to adapt as appropriate.)	<pre>! ap regulatory-domain-profile <profile name> no valid-11g-40mhz-channel-pair 1-5 no valid-11g-40mhz-channel-pair 7-11 no valid-11a-40mhz-channel-pair 36-40 no valid-11a-40mhz-channel-pair 44-48 no valid-11a-40mhz-channel-pair 149-153 no valid-11a-40mhz-channel-pair 44-48 no valid-11a-80mhz-channel-group 36-48 no valid-11a-80mhz-channel-group 149-161 !</pre>

Disabling 40-MHz and 80-MHz Channel Widths

To remove bonded channels from the system and only use 20-MHz channel widths, change the *ARM profile* and the *HT-SSID Profile*. See [Table EC3-1 on page 41](#) for details.

Setting Up a Static 1, 6, 11 Channel Plan for a Stadium

We've come to the most important part of the chapter, where you will learn how to use ARM and *regulatory domain profiles* to construct static channel plans automatically. This technique greatly reduces the workload of the WLAN engineer to configure a large VHD system by combining the best of a manual plan with automated execution.

The idea is to divide the available 2.4-GHz and 5-GHz channels into three different *regulatory domain profiles*. Each profile has a single 2.4-GHz channel – 1, 6, or 11. Evenly distribute the 5-GHz channels across these three groups. For example, [Table EC4-3](#) shows this type of breakdown for an outdoor stadium bowl.

Table EC4-3 Sample Channel Allocation for Stadium Regulatory Domain Profiles

Profile Name	2.4 GHz			5 GHz																									
	ISM			UNII-1				UNII-2				UNII-2 EXTENDED								UNII-3				ISM					
	1	6	11	36	40	44	48	52	56	60	64	100	104	108	112	116	120	124	128	132	136	140	144	149	153	157	161	165	
Stadium-Ch1-Regdomain	✓	--	--	Reserved for Skybox APs				--	--	✓	--	--	✓	--	--	--	Not Allowed in USA				--	--	✓	--	--	✓	--	--	✓
Stadium-Ch6-Regdomain	--	✓	--					✓	--	--	✓	--	--	✓	--	✓					✓	✓	--	--	✓	--	--		
Stadium-Ch11-Regdomain	--	--	✓					--	✓	--	--	✓	--	--	✓	--					--	✓	--	--	✓	--	✓	--	✓

Note that we have reserved four channels for skybox APs to minimize interference with the bowl seating. Also, channel 144 is not used. A total of 17 channels are available with this plan for the bowl seating area.



Channels 120-128 will be allowed in the US again on new APs that pass FCC certification in 2015 and beyond.

Creating the Profiles

To translate this plan into a controller configuration, create a new group and put only those channels into it. (Remove the 40-MHz and 80-MHz channel pairs as well.) [Figure EC4-4](#) is a view of the “-1” group in the ArubaOS GUI.

Regulatory Domain profile > Auditorium-Ch1-regdomain Show Reference Save As Reset

Country Code	US - United States
Valid 802.11g channel	☒ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐ 6 ☐ 7 ☐ 8 ☐ 9 ☐ 10 ☐ 11
Valid 802.11a channel	☒ 36 ☐ 40 ☐ 44 ☒ 48 ☐ 52 ☐ 56 ☒ 60 ☐ 64 ☐ 100 ☒ 104 ☐ 108 ☐ 112 ☒ 116 ☐ 132 ☐ 136 ☒ 140 ☐ 144 ☐ 149 ☒ 153 ☐ 157 ☐ 161 ☒ 165

Figure EC4-4 Sample Channel 1 Regulatory Domain Profile

Here are the equivalent CLI commands to adapt the default US *regulatory domain profile* into a "-1" group that matches [Table EC4-3](#):

Remove channels 6 and 11. (Depending on your country, these may already be enabled by default.)	! ap regulatory-domain-profile <profile name> no valid-11g-channel 6 no valid-11g-channel 11
Remove extra default 5-GHz channels.	no valid-11a-channel 36 no valid-11a-channel 40 no valid-11a-channel 44 no valid-11a-channel 48 no valid-11a-channel 149 no valid-11a-channel 157 no valid-11a-channel 161
Add relevant 5-GHz DFS channels.	valid-11a-channel 60 valid-11a-channel 104 valid-11a-channel 140 !

Repeat the procedure above to create the -6 and -11 groups. Begin each one by creating a new profile. Then adapt the specific channels that are added and deleted based on [Table EC4-3](#). Remember that different country codes begin with different default values, so be sure to adapt accordingly.

Assigning Profiles to the AP Groups

Every AP group has only one associated *regulatory domain profile*. If you have already created the -1, -6, and -11 AP groups, the next step is to assign the appropriate new *regulatory domain profile* to its AP group. [Figure EC4-5](#) shows how to do this in the GUI.




Figure EC4-5 Assigning "-1" Regulatory Domain Profile to "-1" AP Group

Here are the equivalent CLI commands:

Assign regulatory domain profiles to AP Groups.

```
!
ap-group <channel 1 group name>
    regulatory-domain-profile <channel 1 profile>
ap-group <channel 6 group name>
    regulatory-domain-profile <channel 6 profile>
ap-group <channel 11 group name>
    regulatory-domain-profile <channel 11 profile>
```

Provisioning APs Into AP Groups

The final step is to place APs into the correct group. You already should be familiar with provisioning Aruba APs. We strongly recommend that you use the GUI for this step.

To complete this step, you need:

- A CAD drawing or other schematic map showing the APs locations with a repeating 1, 6, 11 layout
- The name of each AP either written on the drawing, or the drawing coded to refer to a separate list / spreadsheet
- A table of MAC addresses mapped to each AP name
- All of the APs to be booted up and visible in the controller

This step can take a long time, depending on the total number of APs in the deployment. You must complete this step anyway, no matter how you decide to plan your channels. Every AP must be named and provisioned before it can begin servicing users.

The incremental work during provisioning beyond what you may already be used to with an Aruba system is to assign each AP to its -1, -6, or -11 AP group. You must flip back and forth to the map quite a bit. But when you are finished, the controller will do all the work from that point forward to produce a fully randomized 5-GHz channel plan based on your static 2.4-GHz assignments.

Setting Up Other Local Channel Plans

The custom *regulatory domain profile* technique can be used for any discrete area where you want to offer a limited set of channels, even just a single channel if necessary (such as for a reserved house channel).

For example, notice in [Table EC4-3 on page 72](#) that the lower four channels are reserved for skyboxes, so we cannot use them in the bowl. Instead, we create an additional regulatory domain profile for the skyboxes all around the bowl. This domain eliminates CCI between skyboxes and the bowl, which is a big limit on performance in non-DFS stadium channel plans. The channel assignments for this new skybox-only profile are shown in [Table EC4-4](#).

Table EC4-4 Sample Channel Allocation for Stadium Skyboxes

Profile Name	2.4 GHz			5 GHz																								
	ISM			UNII-1				UNII-2				UNII-2 EXTENDED												UNII-3				ISM
	1	6	11	36	40	44	48	52	56	60	64	100	104	108	112	116	120	124	128	132	136	140	144	149	153	157	161	165
Skyboxes	✓	✓	✓	✓	✓	✓	✓	--	--	--	--	--	--	--	--	--	Not Allowed in USA			--	--	--	--	--	--	--	--	--

Unlike the seating bowl APs, you only need one AP group and one new *regulatory domain profile* to apply this configuration. Then, during provisioning, you gather up all the APs in the skyboxes and provision them into that AP group. ARM automatically assigns these four channels in a logical way.

Skyboxes are interesting because they are basically in a ring around a central sporting / performance area. From an RF point of view, they are in a linear deployment. Unlike the bowl APs, they will be strongly differentiated from one another. So ARM should have no trouble producing a logical plan.

Configuring Repeating and Non-Repeating Channel Plans

The beginning of this chapter stated that a channel plan is repeating if the same channel number is used more than once in the same coverage area. A channel plan is non-repeating if it cannot reuse the same channel number in the same coverage area.

A channel plan is always repeating unless you explicitly configure it to be non-repeating.

Non-repeating plans must also be static, local plans. You cannot rely on an automated algorithm to determine the architect's intent. So, you must use something like the 1, 6, 11 technique but scaled up to the full number of 5-GHz channels you intend to use. Separate AP groups and *regulatory domain profiles* are required for every 5-GHz channel. You then manually provision one AP in each room into its assigned channel.

This approach is primarily of interest in lecture halls and convention centers with multiple, adjacent VHD rooms. It is more work up front to ensure that the channel allocations are unique in each room, but the result is more than worth it for the increased capacity.

Configuring the Power Plan

As stated earlier, we use ARM for channel control so we must also use ARM for power management. Therefore, configuration is done in the *ARM profile* instead of the *dot11radio profile*.

In the *ARM profile*, the *max-TX-power* and *min-TX-power* settings are adjustable in increments of 3 dBm. ARM is free to choose any value in between or equal to these values. VHD areas usually use fixed power levels determined by surveys. In this case, you would use the same value for both min and max.

If you prefer to attempt to define a power range, do not separate the min and max by more than 3 dB. In general, setting a range simply causes power levels to flap in VHD areas which creates ongoing and unnecessary log messages.

See [Table EC3-3 on page 42](#) for settings to control transmit EIRP using ARM.

Chapter EC-5: SSIDs, Authentication, and Security

Previous chapters in this guide described how to estimate available bandwidth, optimize airtime usage, and construct channel plans. This chapter describes the user experience of finding, joining, and using the wireless network in a VHD area.

Overview

A very high-density (VHD) environment requires many unique considerations for guest and secure Service Set Identifiers (SSIDs).

By definition, VHD networks serve large numbers of guest users. These networks cause obvious challenges from the perspectives of network addressing, scalability, and latency. In addition, challenges are seen in the areas of authentication, hardening, content control, and liability.

In recent years, many large VHD networks have been funded by cellular operators to help offload data traffic from their networks. Wireline and cable operators also have been funding VHD networks to obtain a wireless “play” to help compete with the cellular companies. When sponsored or subscriber-based access is involved, additional back-end authentication, authorization, and accounting (AAA) issues must be factored in. The Passpoint certification from the Wi-Fi Alliance, based on the 802.11u amendment also known as Hotspot 2.0, is one method to address these needs. Passpoint has very specific requirements for SSID design.

Last, but certainly not least, many different user communities work in the VHD facility on a daily basis. Unlike temporary visitors, these users need ongoing, secure, encrypted wireless access. Historically, the method of serving these groups has been a proliferation of separate, preshared key (PSK)-encrypted SSIDs. This scattered approach is no longer viable from an airtime conservation perspective. A new approach is needed.

This chapter moves through these topics in turn:



It is assumed that the reader is fully conversant in the concept of user roles as embodied in ArubaOS, as well as authoring of complex session ACLs.

SSID Strategy for VHD Networks

The critical recommendation in this chapter is to throw out the old model of having separate SSIDs for every different user group in a VHD area.

Airtime is an extremely limited resource. To maximize performance, you must use the absolute minimum number of SSIDs. When SSIDs are created for specific user groups such as staff or press or ticketing, they generate significant additional beacon traffic throughout the area, which leaves less airtime for actual user traffic.

SSID Summary

The Aruba recommended vision is illustrated in [Figure EC5-1](#). This vision ensures maximum performance for all users (guests, employees, and temporary workers) while vastly increasing the security beyond the limitations of preshared keys commonly found in most VHD areas today.




Figure EC5-1 Traffic Flow for Various VHD User Communities

[Table EC5-1](#) summarizes this vision in more concrete terms. Implementation of the open and the converged secure SSIDs is described later in this section.

Table EC5-1 SSID Summary for VHD Design with Converged Secure SSID

SSID Name	Roles Served	Encryption	Authentication	Bands
“Open”	Fan / Guest	None (Open Auth)	Captive portal	Both
“Secure”	- Employees - IT Admin - Team Guest - Concession Vendor - Point of sale terminals - Ticketing - Press - Hotspot 2.0	WPA2-Enterprise (AES with 802.1X)	Multiple different AAA backends, proxied through ClearPass. ClearPass local DB can be used to administer accounts as well.	5 GHz only

This strategy integrates cleanly with Hotspot 2.0 / Passpoint deployments, and it avoids the deployment of an additional SSID that would otherwise be required.

Open Guest Network with Captive Portal

The guest network is the centerpiece of any VHD wireless network. It must:

- Use open 802.11 authentication (not require encryption)
- Offer a captive portal with an acceptance of terms of service
- Operate on both bands for compatibility with legacy devices
- Block hacking or malicious activity via layered firewall rules

Converge Secure Users on to One SSID

Use the dynamic role and VLAN assignment capability of the Aruba controller to collapse all of the different secure users onto one SSID. Specifically, Aruba makes these recommendations:

- A single, secure SSID must be broadcast alongside the open SSID for guests
- The secure SSID is accessible only on the 5-GHz band
- The secure SSID require 802.1X authentication
- PSK SSIDs are prohibited in the facility to eliminate unnecessary beacons
- Third-party APs are prohibited to prevent misconfigured SSIDs from reducing available capacity
- A ClearPass server is deployed to serve as a clearinghouse to process authentication requests from all the different user groups. Depending on the user group, ClearPass itself could be the authentication server, or ClearPass could proxy to another AAA server.
- The controller(s) must submit AAA requests to the ClearPass server(s)
- User-derived or server-derived roles must be used to place each user into the correct VLAN

Role and Policy Strategy for VHD Networks

Roles are different from SSIDs.

Aruba knows from deploying many different kinds of VHD networks that virtually every one has a minimum of three roles:

- Guest/student/fan Internet role
- Facility employee role
- Facility IT admin role

In addition, depending on the type of facility, one or more of these roles are also common:

- Concession vendor point-of-sale role
- Ticketing role
- Press / media role
- Carrier offload role
- Performer / entertainer role
- Building controls admin role
- Physical security role
- Police / Fire / EMS role

Each of these roles is associated with a discrete VLAN and subnet. The VLANs must be trunked to the Aruba controllers so that users can reach destinations on those nets.

Table EC5-2 is a summary of these typical user roles common to all VHD facilities, the corresponding use cases, and a high-level list of important network policies.

Table EC5-2 User Role Breakdown for All VHD Facilities

Role	Use Case	User Count	Security	DHCP Source	Session "Permit" Rules	Session "Deny" Rules
Guest / Fan	General public internet access	- Expected users = Guest ADC - Mix of tablets and smartphones	Open SSID with captive portal. The CP typically contains the acceptable use policy and any other legal messaging that is required to protect the facility operator.	High-capacity DHCP server	Mandatory: - Standard captive portal rules - DHCP and DNS lookup - HTTP/HTTPS to Internet - Social media - Skype / Lync / FaceTime - Mobile app stores - VPN to Internet - VPN to local IT jump box for live diagnostics - Speedtest.net Case by case: - QoS remarking of video or other traffic for "house" applications	Mandatory: - Anything that is not permitted - Peer-to-peer (all forms) - Chatty protocols (mDNS, IPv6) - Offering of network services (DHCP, DNS) - Any other non-Internet access or access to private subnets Case by case: - Operating system and application software updates - Time of day restrictions
Facility Employee	- Internal resources using venue-owned devices - Trusted vendors or contractors using venue-owned devices and applications	- Expected users = Staff ADC - Mix of laptop, tablet and smartphones	802.1X	Conventional DHCP server	Match existing network policy (or leverage upstream firewall)	Match existing network policy (or leverage upstream firewall)
Facility IT Administrator	Facility network administrators	- Less than 50 - Mix of laptop, tablet, and smartphones	802.1X	Conventional DHCP server	Everything	Nothing

Table EC5-3 Additional User Role Breakdown for Stadiums, Arenas, Airports, and Convention Centers

Role	Use Case	User Count	Security	DHCP Source	Session "Permit" Rules	Session "Deny" Rules
Concession Point-of-Sale	- Mobile point-of-sale devices used by concession vendor - Accessible stadium-wide	- Less than 512 - Could be modern tablets, phones running an app, or legacy hardened scanning terminals	802.1X	POS Firewall (typically also with MAC address enforcement of DHCP leases)	Usually implemented on vendor firewall	Usually implemented on vendor firewall
Ticket scanners	- Mobile ticket scanning terminals - Need to connect to ticket vendor servers, the venue back office application, or both - Accessible in the parking lots, stadium perimeter, and all building entrances	- Less than 150 - Could be modern tablets, phones running an app, or legacy hardened scanning terminals	802.1X	Ticket vendor firewall (with MAC address enforcement of DHCP leases)	Usually implemented on vendor firewall	Usually implemented on vendor firewall
Press	- Dedicated SSID with no rate limits and high QoS for press users, especially photographers - Limit press SSID to only those APs that require it - Accessible on the field/court, in the press box, and possibly in dedicated press areas by locker rooms	- Less than 250 - Mix of laptop, tablet, and smartphones - Usually very old devices with poor drivers	802.1X or PSK (802.1X preferred)	Conventional DHCP server	Same as fan role with QoS marking	Same as fan role
Police / Fire / EMS	Mobile data terminals used by public safety personnel	- Less than 50 - Mix of laptop, tablet, and smartphones	802.1X	Agency firewall (with MAC address enforcement for DHCP leases)	Usually implemented on agency firewall	Usually implemented on agency firewall

Guest SSID Design

The open SSID for guests automatically redirects users to a standard captive portal for them to agree to the acceptable use policy (AUP). In addition, the configuration of the SSID itself needs certain options enabled to support very high user loads.

Captive Portal

Aruba offers a full-featured, highly available captive portal solution called ClearPass Guest. This solution allows for customizable skins that adapt to various screen sizes. Delivery of advertisements is supported, as is any level of data collection desired by the facility operator. Full RADIUS accounting and RFC3576 COA are supported. Depending on the number of users, Aruba generally recommends deploying two servers in a cluster configuration.

The screen shots in [Figure EC5-2](#) show these capabilities:

- Station Casinos: Simple one-click button to acknowledge AUP and gain Internet access
- Nationwide Arena: Gathers an email address, delivers a banner ad, and includes AUP acknowledgment
- Levi's Stadium: Gathers name and email address, plus AUP acceptance




Figure EC5-2 ClearPass Guest Skins

For liability reasons, Aruba recommends you always use an AUP and deploy a captive portal solution. If you have an existing captive portal appliance, the Aruba controller can be configured to use it.

See the *ClearPass User Guide*, which contains detailed material on how to design and deploy the solution. The controller requires special configuration to perform user redirection and COA enforcement.



The built-in controller captive portal is designed for low-transaction rate environments, and it should not be used for VHD areas.

Certificates

You must obtain and install certificates from a valid certificate authority (CA) when you run a captive portal. A server certificate installed in the controller verifies the authenticity of the controller for captive portal. Without a valid certificate, all guest users who go through the captive portal get an error message about the certificate.

Aruba controllers ship with a demonstration digital certificate. Until you install a customer-specific server certificate in the controller, this demonstration certificate is used by default for all secure HTTP connections such as captive portal. This certificate is included primarily for the purposes of feature demonstration and convenience. It is not intended for long-term use in production networks. Users in a production environment are urged to obtain and install a certificate issued for their site or domain by a well-known CA. You can generate a certificate signing request (CSR) on the controller to submit to a CA.

MAC Caching

If the guest disconnects and reconnects while the account is valid, the device MAC address is used to re-authenticate the user. The guest user need not re-enter credentials after a successful initial authentication.

SSID and Virtual AP Configuration

The configuration of the guest SSID must include a number of changes to the default settings. These changes repeat certain configuration items discussed in [Chapter EC-3: Airtime Management](#).

Table EC5-4 SSID Related Configuration for Guest SSID

Profile	Option	Setting	Reason
Virtual AP	deny-inter-user-traffic	Enabled	Prevents peer-to-peer communication.
Virtual AP	broadcast-filter-all	Enabled	Prevents broadcast and multicast traffic from being sent over the air. (Unless multicast is required for a specific reason – not recommended in VHD areas.)
SSID	max-clients	250	Default value of 64 is too small for VHD areas.
Traffic Management	shaping-policy	preferred-access	ATF rate limits the air and is vital to prevent older, slower clients from starving newer, faster clients.

Secure SSID Design

The secure SSID should be implemented with a standard WPA2-Enterprise configuration. ClearPass Policy Manager (CPPM) should be deployed as the AAA server for the secure network. CPPM can proxy to other AAA sources as necessary to authenticate roles that are not served from the local database. CPPM supports Active Directory, LDAP, RADIUS, and other major AAA integration protocols.

Depending on the number of users and how mission-critical the secure SSID is, Aruba generally recommends deploying CPPM in a cluster configuration. The Aruba controller should be configured to load balance requests across the cluster.

Compatibility Requirements

To achieve the vision of a single, converged, secure SSID, you must work through a diligence process with each of the user communities who exist in the coverage area:

- To ensure the devices used support 5-GHz operation (and DFS channels)
- To ensure the devices support WPA2-Enterprise with AES encryption
- To ensure the devices used have 802.1X supplicants
- To test that the supplicants are compatible with the system
- To identify the most appropriate AAA repository for each group, and how it will be queried over the network
- To identify the best EAP type for the 802.1X authentication
- To develop a plan to phase out PSKs and 2.4-GHz usage, including replacement or upgrading of devices as necessary

Aruba has had good success working through this process with major ticketing, point-of-sale, and other vendor types to implement this design.

Preshared Key Security Risks

Some user groups with PSKs may not give them up easily. PSKs are preprogrammed into a device, which often eliminates the need for individual employees to log into the network, and facilitates sharing of devices. This simplicity is highly desirable for high-turnover roles like concession workers, or jobs where the same devices change hands many times in the course of the same shift.

The resistance usually has two causes:

- The actual or perceived work needed to reconfigure their devices
- The cost to upgrade devices that are found to be incompatible

One of your response strategies for forcing the move to 802.1X is the many inherent security risks of PSKs. If you are providing the service to all the vendors, then these risks directly expose the entire system:

- **Theft:** Mobile devices are sometimes stolen for cash, along with the security credentials. However, a hacker could steal a device specifically to recover the PSK that might be displayed in cleartext on older devices. Or a hacker could use a device with a valid PSK to gain access to the vendor network to perform other attacks.
- **Hackable:** As computers get faster and prices drop, older encryption technologies have become vulnerable. Recently, weaknesses in WPA-PSK encryption have been identified, and it is no longer considered secure. WPA2-PSK is currently considered secure, but this will inevitably change.
- **Untraceable:** PSKs, by definition, are the same on every device. It can be nearly impossible to trace a security incident to a given person.
- **Static:** PSKs are almost never changed, especially on older terminals with slow user interfaces. Hackers have more incentive to crack a password because it can be used for a long time.
- **Hard to Change:** For the same reason, even after a known security breach or device theft, most PSKs are never changed, even though the organization may be aware they face some risk.

PSKs therefore pose a threat to the security of modern Wi-Fi networks.

Replicating the PSK Experience with 802.1X

There is a solution if you have a user community that would never normally assign per-user credentials to their devices, and legitimately needs the simplicity of the PSK.

The PSK experience can be replicated easily under 802.1X through the use of a static username+password combination. For example, all your ticket scanners could have username “ticket” and password “changeme”. This method is functionally equivalent to a Wi-Fi PSK, and it allows all of these devices to coexist within the 802.1X security framework.

This method is more secure than PSKs because while the credentials are still statically assigned, they are protected inside the EAP tunnel in the 802.1X exchange.

Aruba further recommends adding MAC address enforcement for any such devices. ClearPass Policy Manager can be configured with a local database of valid MAC address. The authentication rule to grant access can require that:

- The username+password be correct
- The MAC address be valid

This method allows you to potentially “kill” a device that goes missing without compromising all of the rest of the devices.

Zero Tolerance for Third-Party APs

An important part of the converged secure SSID vision is that 100% of vendors that need Wi-Fi must use the house network. No vendor may install or operate its own radio equipment, including mobile hotspots using 3G/4G backhaul that advertise Wi-Fi access. Speakers, sporting teams, or performers who come onsite for a few days cannot bring their own IT setup. The house network must be used to minimize intersystem interference and resulting performance impacts for all users.

Aruba experience shows that third-party APs are rarely if ever tuned for high-density operation. Specifically, they use low data rates, low beacon rates, and they violate most of the best practices in this VRD. Most third-party APs use 1-Mbps beacon rates on the 2.4-GHz band and 6 Mbps on the 5-GHz band, which will negatively affect all users on these bands. The only way to avoid these issues is to forbid the use of these APs.

When vendors are placed onto the house system, which has been fully tuned for VHD operation, they will experience improved connectivity. Everyone has an incentive to participate.

You must set a zero-tolerance policy for this strategy to work. Third-party radio systems should be detected and located by the facility operator and removed at the vendor’s expense.



The only viable alternative is to set aside one or two “house” channels for exclusive use by third-party APs. See [Exception #3 – “House” Channels on page 24](#) for more discussion. However, the venue then must communicate, monitor, and enforce the use of house channels by short-term guests with their own IT equipment.

SSID and Virtual AP Configuration on Controller

The configuration of the secure SSID must include changes to the default settings (Table EC5-5). These changes are in addition to the data rate and other configuration items discussed in [Chapter EC-3: Airtime Management](#).

Table EC5-5 SSID-Related Configuration for Secure 802.1X SSID

Profile	Option	Setting	Reason
Virtual AP	broadcast-filter-all	Enabled	Prevents broadcast and multicast traffic from being sent over the air. (Unless multicast is required for a specific reason – not recommended in VHD areas.)
SSID	opmode	WPA2-Enterprise	For multiple user groups to share a single secure SSID, it must use 802.1X.
AAA Authentication Server Group	set role	condition Class value-of	Enables server-derivation rules.
	set vlan	condition Aruba-User-Role contains <CPPM> role set-value <VLAN>	Uses one entry for each role+VLAN combination that could be returned by CPPM.
	load-balance	Enabled	Distributes requests across AAA cluster.
Traffic Management	shaping-policy	preferred-access	ATF rate limits the air and is vital to prevent older, slower clients from starving newer, faster clients.



It is assumed the reader is familiar with RADIUS server configuration on the Aruba controller. This configuration does not include every setting required for WPA2-Enterprise SSID.

Dynamic Authorization Profiles on ClearPass Policy Manager

ClearPass manages staff and vendor devices for wireless access to the VHD network. Authentication requests from staff devices are typically proxied to an Active Directory or other enterprise server to which ClearPass has been securely joined.

The ClearPass local user database can be used for other user groups such as ticket scanners or one-off contractors who would not normally receive accounts in the enterprise AAA server. The local database would also serve staff and vendors that only authenticate using a nonregistered or domain device that is redirected to the ClearPass onboard captive portal page and completes the onboarding process to securely access the network.




Figure EC5-3 Authentication Workflow for Secure SSID

An 802.1X service on ClearPass is used to receive the RADIUS request from the controllers, and sequentially query defined AAA sources to validate the supplied credentials. [Figure EC5-4](#) shows a CPPM 802.1X service definition for seven different typical secure roles found in a stadium/arena type environment. This screenshot is meant only to illustrate the capability. The exact policy configuration should be developed uniquely for each VHD facility.




Figure EC5-4 ClearPass 802.1X Service Definition for Multi-Role Derivation

MAC Address Authentication

To use MAC address authentication in addition to either static or dynamic username/password authentication, use ClearPass. In this case, the enforcement policy on ClearPass is configured to grant access to the user only when user ID, password, and device MAC address all match information stored in ClearPass.

In the example in [Figure EC5-4](#), rows 4 and 5 show two roles configured in this way. The “Tips:Role” must equal a valid system role (which implies that the username and password have been validated). The “Tips:Role” must also match a Valid Vendor Device. This rule implements MAC authentication.

ClearPass offers various solutions for managing databases of MAC addresses in a simple way.

ClearPass Policy Design

Actual configuration of ClearPass is outside the scope of this VRD. Consult with your local Aruba systems engineer or with a ClearPass-certified Aruba integrator.

The purpose of going into some detail on the security architecture is to describe the full concept and the roadmap to implementing it. Readers who have worked only in PSK environments may not otherwise think to go this direction. However, the airtime gains that are possible with this approach are so substantial in VHD environments as to warrant investigation. All VHD areas have at least a few user groups that require secure access, so the strategy can be leveraged by all readers.

Offload SSID Design

Three major technologies are used for service providers to deliver Wi-Fi services to their subscribers:

- Passpoint is the next-generation solution ratified and certified by the Wi-Fi Alliance for EAP-SIM, EAP-TTLS, and EAP-AKA on 802.1X-enabled SSIDs.
- Wireless Internet Service Provider Roaming (WISPR) is an older technology that some operators still use.
- Native EAP-SIM and EAP-AKA on dedicated 802.1X-enabled SSIDs can be enabled without employing Passpoint.

Aruba supports all three solutions. In general, the market is moving away from WISPR and native EAP protocols for a variety of reasons. Consult with your Aruba systems engineer if you need to use either in your VHD environment.

Passpoint

ArubaOS incorporates Passpoint technology from the WFA Hotspot 2.0 Specification to simplify and automate access to public Wi-Fi networks. Service provider Wi-Fi networks increasingly use Passpoint for authentication.

Hotspot 2.0 is based upon the 802.11u protocol. This protocol provides wireless clients with a streamlined mechanism to discover and authenticate to suitable networks. The 802.11u protocol allows mobile users to roam between partner networks without additional authentication.

The SSID architecture envisioned by Passpoint aligns extremely well with the recommendations in this chapter. Passpoint envisions two different SSIDs used in most public spaces:

- An open SSID used solely for registration of new subscribers
- An encrypted 802.1X SSID for all subscriber communications

Subscribers who have previously completed a registration and have an 802.1X supplicant already configured automatically connect to a Passpoint-enabled SSID that advertises a compatible operator network ID.

Passpoint configuration and architecture is beyond the scope of this VRD. See the Aruba white paper “Wi-Fi Certified Passpoint Architecture for Public Access.”

Personal Hotspots

So-called “My-Fi” devices are an increasingly common feature everywhere, especially high-density environments. Smartphone users may intentionally or unintentionally have tethering services enabled. Numerous other Wi-Fi enabled devices have Basic Service Set (BSS) or Independent Basic Service Set (IBSS) functionality, such as GoPro cameras, which are common in sporting events. Fortunately, at the time of this writing, virtually all My-Fi devices use the 2.4-GHz band for operation. Unfortunately, they use 1-Mbps beacon rates and other inefficient configuration options.

Unless you want to search everyone at the door and confiscate such devices, nothing realistically can be done about these devices.

Do not be tempted to use wireless intrusion protection system (WIPS) containment on these devices. Containment magnifies the problem by creating even more control traffic in a fruitless attempt to squelch these sources.

The best solution is to focus on moving every possible user to the 5-GHz band. Use the RF management practices discussed in [Chapter EC-3: Airtime Management](#) to increase the percentage of 5-GHz users.

Post-Deployment Site Surveys

RF site surveys are mandatory for large VHD facilities like airports, arenas, convention centers, and stadiums after the APs are installed. A survey is the only way to establish the proper EIRP level for the APs. You may need to run multiple surveys at different power levels to determine in the correct level.

Aruba supports the leading site survey solutions from AirMagnet and Ekahau. To simplify your work, enable the broadcasting of AP names by the access point. When enabled, this feature adds a special IE to the beacon that can be read by these tools to display the AP name as provisioned on the controller. Disable this feature when you finish the survey to reduce the beacon size.

As mentioned in [Chapter EC-2: Estimating System Throughput](#), it is critical that you survey with 1SS phone clients. Ekahau and AirMagnet offer Android-based survey utilities that integrate with their PC-based platforms. Aruba has observed a 10 - 20 dB signal difference between phones and laptops when performing site surveys.

Table EC5-6 Profiles and Settings for Site Surveying

Profile	Option	Setting	Reason
SSID Profile	advertise-ap-name	Enable	Causes the APs to add an information element that includes the name of each AP in its beacons. Should be disabled when surveys are completed.

VHD Network Hardening

Open guest networks in VHD areas are natural targets for casual and malevolent hackers. Aruba recommends these hardening best practices:

- If possible, do not configure a Layer 3 interface on wireless user subnets (including secure subnets) unless a captive portal is being used and redirect is required. The controller should be Layer 2 only on all wireless subnets into which users can be placed.
- Do not be the default gateway for any user subnet.
- Place the DHCP server on a Layer 3 separated subnet and use a helper.
- Do not be the IP helper for any user subnet.
- Configure Aruba Validuser ACL to allow specific user subnets, and disallow protected host IPs or IP ranges including default gateways for each VLAN, DNS, DHCP, etc.
- The guest role should explicitly disallow connection to network infrastructure elements via TCP ports 22 and 4343.
- The guest role should explicitly disallow Telnet, SSH, and other protocols that are not required for guest services.
- Enable ARP spoof prevention on the default gateway for wireless user subnets, and also on the controller if it has an L3 address on any wireless user VLANs.
- Use a dedicated infrastructure subnet to connect all Wi-Fi controllers, APs, and servers.
- Use ClearPass for administrator authentication using RADIUS and/or TACACS. Monitor the logs.
- Use an IDS solution to monitor infrastructure and user subnets for suspicious activity.
- Shut all unused Ethernet interfaces on the controller.
- Monitor for rogue and potential rogue devices in AirWave and on the controller.
- Enable “enforce-dhcp” in AAA profiles to prevent users from being able to assign static addresses and gain access to disallowed networks or spoof network resources.

Links to ClearPass Technical Documents

Table EC5-7 ClearPass Configuration and User Guides

Document Title	Author	Link
ClearPass Policy Manager 6.3 User Guide	Aruba Networks	http://community.arubanetworks.com/t5/Software-User-Reference-Guides/ClearPass-Policy-Manager-6-3-User-Guide/ta-p/168394
ClearPass Guest 6.4 User Guide	Aruba Networks	http://community.arubanetworks.com/aruba/attachments/aruba/SoftwareUserReferenceGuides/29/1/ClearPass_Guest_6.4_User_Guide.pdf
ClearPass Clustering Design Guidelines	Aruba Networks	https://arubapedia.arubanetworks.com/arubapedia/images/4/43/CPPM_TechNote_-_Clustering_Design_Guidelines_V1.pdf
ClearPass Load Balancing and F5 BIG-IP Integration	Aruba Networks	https://arubapedia.arubanetworks.com/arubapedia/images/8/83/CPPM_Load-Balancing_TechNote_v1.0.pdf

Chapter EC-6: Video Streaming

Video streaming in very high-density (VHD) environments is a growing area and is one of the more exciting applications. Video streaming is also one of the most technically difficult applications to deliver successfully. Video is unlike most other network loads in that each stream has a nearly constant duty cycle and requires high bandwidth. This chapter considers how to dimension and deploy a video service in VHD facilities.

Dimensioning Video Usage

For success with video, you must set proper expectations. Spectrum is a finite resource and VHD areas tend to be single collision domains with little or no RF spatial reuse. As a result, there are distinct limitations to how much traffic can be carried by any VHD system. These limits are true regardless of the WLAN vendor, as well as whether unicast or multicast are used.

Dimensioning Process

The basic flow is shown in the flowchart. In general, the dimensioning process determines the allowable bitrate of the video stream. This is the opposite of how wired video distribution systems are dimensioned, which can scale with offered load. In the case of Wi-Fi, the spectral capacity is fixed. Therefore, the video bitrate is conditional on the TST formula and the number of expected video users.



Begin with the Guest ADC. You learned to compute associated device capacity (ADC) as part of system dimensioning in [Chapter P-2: System Dimensioning](#) of the *Very High-Density 802.11ac Networks Planning Guide*. Step two is to compute the expected TST for the facility. TST was described in [Chapter EC-2: Estimating System Throughput](#) of this guide. The TST concept is important when you consider video load.

Next you estimate the take rate. Fortunately, video usage levels are generally a small percentage of seating capacity. Plan no more than 50% maximum concurrent video usage. This percentage can vary widely in different types of facilities:

- Lecture halls with e-learning applications may range between 20-50%
- Airports with passengers streaming movies may be a little lower
- Stadiums and arenas tend to be 10% or less as of the date of writing

These three VHD video use cases are probably the most dominant. If video is not an actual service that is offered by the facility, it need not be explicitly dimensioned and can be treated as part of the overall TST analysis.

The next step is to choose the usage mode: either live streaming or video on demand (VOD). This matters because live streaming services can potentially leverage multicast. VOD systems by definition require unicast flows to each user. If your streaming solution allows each user to start a prerecorded stream on demand (instead of coming in the middle of an existing live stream) then it is VOD.

You may be surprised to learn that true video streaming does not necessarily save any bandwidth at the end of the day. In fact it can consume significantly more bandwidth than unicast! We'll cover this in a moment.



This VRD is not concerned with two-way videoconferencing, which is not a typical VHD application.

The last step is to calculate the available maximum bitrate that the system can support based on the previous selections.

Computing a Bitrate Table for a Stadium

Due to the very wide range in video take rates experienced in many VHD facilities, Aruba recommends that you compute a table of potential bitrates based on various scenarios. Such a table also allows you to perform “what if” analysis for RF spatial reuse if you have evidence that your specific system is capable of reuse.

Consider [Example #3 – Outdoor Stadium in USA on page 33 in Chapter EC-2: Estimating System Throughput](#). In this example, we have:

- An ADC of 30,000
- A band split of 75% on 5-GHz and 25% on 2.4-GHz
- A 5-GHz TST of 630 Mbps with 1 reuse
- A 2.4-GHz TST of 60 Mbps with 1 reuse

With this information, we can construct a table of possibilities for 5-GHz by simply dividing the user count for various video take rates into the TST.

Table EC6-1 Example Video Bitrate Scenarios for Outdoor Stadium (5-GHz)

Take Rate	Concurrent	5 GHz			
		1 Reuses	2 Reuses	3 Reuses	4 Reuses
	Video Users	630 Mbps	1,260 Mbps	1,890 Mbps	2,520 Mbps
5%	1,125	560 Kbps	1,120 Kbps	1,680 Kbps	2,240 Kbps
10%	2,250	280 Kbps	560 Kbps	840 Kbps	1,120 Kbps
15%	3,375	187 Kbps	373 Kbps	560 Kbps	747 Kbps
20%	4,500	140 Kbps	280 Kbps	420 Kbps	560 Kbps
25%	5,625	112 Kbps	224 Kbps	336 Kbps	448 Kbps
30%	6,750	93 Kbps	187 Kbps	280 Kbps	373 Kbps

These numbers are very sobering. Note that these are gross maximum bitrates, including all video application overhead. Furthermore, we have used 100% of the TST for video, which means that no capacity at all is left for other users. Obviously this is unacceptable.

Table EC6-1 forces the wireless architect to do one or all of these actions:

- Attempt a design that achieved at least two RF spatial reuses.
- Use a very low video bitrate (probably no more than 256 Kbps).
- Obtain a much more precise estimate of the likely video take rate.

We start with this example intentionally to show how difficult video is to achieve in VHD facilities. It gets even worse when we consider the equivalent math for the 2.4-GHz side of the system in [Example #3 – Outdoor Stadium in USA on page 33](#).

Table EC6-2 Example Video Bitrate Scenarios for Outdoor Stadium (2.4-GHz)

Take Rate	Concurrent	2.4 GHz			
		1 Reuses	2 Reuses	3 Reuses	4 Reuses
	Video Users	60 Mbps	120 Mbps	180 Mbps	240 Mbps
5%	375	160 Kbps	320 Kbps	480 Kbps	640 Kbps
10%	750	80 Kbps	160 Kbps	240 Kbps	320 Kbps
15%	1,125	53 Kbps	107 Kbps	160 Kbps	213 Kbps
20%	1,500	40 Kbps	80 Kbps	120 Kbps	160 Kbps
25%	1,875	32 Kbps	64 Kbps	96 Kbps	128 Kbps
30%	2,250	27 Kbps	53 Kbps	80 Kbps	107 Kbps

In effect, [Table EC6-2](#) says that 2.4-GHz is virtually unusable for video even with significant, proven RF spatial reuse. And this table assumes that fully 75% of Wi-Fi users are successfully steered to the 5-GHz band!

Computing a Bitrate Table for a Lecture Hall

Moving to a somewhat easier example, now consider [Example #1 – Small Auditorium on page 31](#). To recap, this example had these attributes:

- An ADC of 1,050
- A band split of 50% on 5-GHz and 50% on 2.4-GHz
- A 5-GHz TST of 284 Mbps
- A 2.4-GHz TST of 201 Mbps

No RF spatial reuse is possible because of the indoor channel model and small physical volume. From these values, you can easily produce [Table EC6-3](#) for different take rates. We start at 20% because it is expected to have a higher take rate in e-learning classrooms.

Table EC6-3 Example Video Bitrate Scenarios for Lecture Hall Example

Take Rate	Concurrent	5 GHz	2.4 GHz
		1 Reuse	1 Reuse
	Video Users	284 Mbps	201 Mbps
20%	105	2,705 Kbps	1,914 Kbps
25%	131	2,164 Kbps	1,531 Kbps
30%	158	1,803 Kbps	1,276 Kbps
35%	184	1,546 Kbps	1,094 Kbps
40%	210	1,352 Kbps	957 Kbps
45%	236	1,202 Kbps	851 Kbps
50%	263	1,079 Kbps	766 Kbps

While these numbers are better than the stadium scenario, some of the same caveats apply. These gross bitrates assume that all system capacity is used for video. For example, if you want to keep video to 50% of total channel load, you would have to cut all of these numbers in half.

A further complication in the adjacent lecture hall scenario is co-channel interference from nearby rooms with same-channel APs. This complication is explained at length with illustrations in the *Very High-Density 802.11ac Networks Scenario 1: Large Adjacent Auditoriums* guide. In the stadium case, the CCI is in effect captured by the low overall TST values. However, the small auditorium example that has just a few channels, no provision is included for this possible impairment.

Multicast vs. Unicast

How does using multicast affect the analysis we just completed? Unfortunately, the answer is “not much.”

As stated earlier, many of the most desirable video applications in large VHD areas like stadiums are instant replay. By definition, replay services are VOD and therefore cannot be delivered with multicast. Even a movie streaming service inside an airplane is VOD because even though all the content is technically streaming, each user can start/stop his or her own stream.

The only candidates for multicast are true live streams, where the viewer joins a program that is already in progress. The stream cannot be paused, or have very limited pause capability. These live streams do exist in some stadiums, for example, an American football stadium that broadcasts the NFL RedZone television channel inside the stadium over Wi-Fi.

Multicast in large scale video systems has four major problems:

- Low multicast data rate
- No acknowledgments
- Multicast replication explosion
- Prevents use of broadcast-multicast filter

Low Multicast Data Rate

In [Chapter EC-2: Estimating System Throughput](#), you learned that multicast frames can be sent only at legacy 802.11a OFDM data rates (from 6 Mbps to a maximum of 54 Mbps). Multicast cannot be sent at modern MCS rates, nor can multiple spatial streams be used. Therefore, a 2SS smartphone capable of a 173.3 Mbps unicast rate would downshift to as little as 6 Mbps to receive a multicast stream.

You also learned that the default behavior in 802.11 is to transmit multicast traffic at the lowest configured basic rate for the BSS, so it stands the best chance of reaching all associated clients. Using the Aruba multicast rate optimization (MRO) feature, this can be increased slightly to the highest **common** low data rate supported by the clients joined to a multicast group. But while this might increase the data rate from the default of 6 Mbps to perhaps 24 Mbps, it is unlikely to take it much higher in practice.

Aruba also supports setting a fixed multicast rate just for video traffic. Configuration of this feature is explained in a few pages.

Another way to affect this problem is to use the Aruba dynamic multicast optimization (DMO) feature to convert multicast streams into unicast streams in real time. This feature generally works well for between 40 and 80 video users per AP, depending on whether they are 802.11n or 802.11ac and how many spatial streams they support. Above the DMO threshold, the system reverts to multicast. But if you are expecting more users than the DMO threshold, then its not likely the system will work anyway.

No Acknowledgments

Per the 802.11 standard, multicast frames are not acknowledged. Lost frames are not detected and not resent. This fact means that multicast video is vulnerable to lossy environments, which is an excellent description of all VHD networks.

By using the Aruba DMO feature to convert multicast to unicast, you also address this problem because unicasts are acknowledged. However, the acknowledgment process itself imposes additional overhead on the channel.

Multicast Replication Explosion

Next, we have the replication explosion problem. This problem affects VHD deployments that have more APs than channels.

Assume you have 20 APs on each channel, and that most of them can hear one another. Further assume that each access point (AP) has just one multicast subscriber. This situation means that each of the 20 APs has to broadcast the exact same multicast stream onto the same channel at the same time.

This aspect of multicast video operation is very important but often is overlooked. Many customers are unaware of this problem and they invest in multicast video systems and plan to scale to large user counts.

Not only is it an enormous waste of bandwidth, but when one considers the points about low data rates, it makes multicast virtually impossible to use in modern Wi-Fi systems with modern smartphones that boast 1080P or better displays. For this reason alone, unicast is generally a superior video delivery method because many more video streams can be carried at higher data rates when there are many same-channel subscribers on different APs.

Prevents Use of Broadcast-Multicast Filter

The best practice VHD configuration shown in [Chapter EC-3: Airtime Management](#) includes enabling the Aruba *broadcast-filter-all* parameter. This hard filter at the radio edge blocks virtually all broadcast and multicast traffic from being transmitted on the air.

This setting is essential to prevent normal network “chatty” traffic such as Layer 2 service discovery protocols and IPv6 multicasts from being forwarded automatically onto the air. When this occurs, precious capacity is consumed that can never be recovered.

To enable multicast video, you must disable this filter, thereby allowing all of these unnecessary transmissions to make it onto the air. Unicast video runs with this filter enabled.

Video Quality vs. System Capacity

With any wireless network, there is an explicit tradeoff between video quality and overall system capacity. As you can see from the tables earlier in this chapter, this tradeoff can be painful in large VHD environments.

The latest generation of smartphones has further complicated this situation due to the arrival of 1080P-equivalent displays. More pixels demand higher bitrates to present a clearer picture. However, the small form factor of a phone screen, no matter what the resolution, tends to make the lower resolution of lower bitrates such as 256 Kbps less noticeable. Note that motion blur for action sports affects all bitrates and tends to equalize visual quality.

However, none of those facts change the basic physics problem of limited RF spatial reuse in VHD environments. The wireless architect must work with the customer to explain the issues involved and to push aggressively for the lowest possible video bitrate.

Aruba recommends this approach:

- Explain the tradeoff between quality and capacity with the project stakeholders responsible for the video system.
- Use absolute minimum bitrate possible for the first few events in order maximize overall system capacity while you evaluate the actual demand level. 256 Kbps is recommended at this stage.
- Study the overall system utilization to determine how much headroom you may have in the TST to accommodate higher bitrates.
- If the video take rate is low enough, and the overall total traffic volume is low enough, experiment by increasing the video bitrate in increments of 256 Kbps.
- 768 Kbps or 1 Mbps are the maximum bitrates that should ever be used in an arena or stadium environment.
- 1.5 Mbps or 2 Mbps are the maximum bitrates that should be used in a lecture hall environment.

Note that the quality of the video encoding hardware and how it is configured has an enormous effect in the resulting video quality on the device. A poorly encoded stream will always look bad no matter what the bitrate, while a well-encoded stream can look exceptional even at 256 Kbps.

Configuring Unicast Video Streaming

In VHD environments with over 1,000 seats, Aruba strongly recommends unicast video as the default choice for VHD environments that offer video streaming. If the video solution offers only multicast, then use DMO to convert the multicast to unicast in realtime.

Unicast is required for all downstream video use cases except for true live streams.

Unicast video requires relatively little additional configuration beyond the best practices listed in [Chapter EC-3: Airtime Management](#). Use the entire configuration from that chapter, being sure to enable *broadcast-filter-all*. Be sure to set your DSCP-to-WMM QoS markings as explained in that chapter.

Ensure End-to-End QoS Marking

As always with any kind of network video, configure all network elements from the video servers all the way to the APs to properly mark video traffic. QoS is critical to ensuring that video traffic receives the proper transmit priority on the air.

Use a packet capture tool to verify QoS on each major traffic egress point in the wired network. Use a wireless packet capture tool to verify that the WMM markings are correct in the MAC header of video frames.

If the markings are not present, troubleshoot the problem and correct the misconfigured network elements.

Configure Remarking ACLs

If your video head end does not mark frames, implement remarking access control lists (ACLs) in the wired network on the switch to which the servers connect.

Remarking can also be done on the Aruba controller, however, by this point the video has already transited one or more switches inside the network and could already be impacted. Remarking should be done at the ingress point of the video traffic.

Verify MTU Size

Check your video maximum transmission unit (MTU) size to ensure that it will traverse the end-to-end network without fragmentation. 1,400 bytes is a conservative MTU value.

Conversely, consider enabling jumbo frames to take advantage of the larger A-MSDU capability of 802.11ac. Aruba APs support A-MSDU values of up to 3, meaning that the AP send individual MSDUs of up to 4,500 bytes. When combined with A-MPDU, this potentially allows the system to significantly reduce the airtime required for each client. A-MSDU applies only to unicast video.

Ensure Video-Aware ARM is Enabled

ARM does not perform periodic off-channel scanning when video users are active on the system. This feature is called video-aware ARM. It is enabled by default in 6.4.2 and higher. If you are running earlier code, verify this setting in the ARM profile. In general, for VHD areas to get the latest performance optimizations, run at least 6.4.2 code.

Configuring Multicast Video Streaming

Multicast may be appropriate for VHD environments of 1,000 seats or less. This threshold is arbitrary and chosen to highlight the distinction between the two methods. Ultimately the wireless architect is responsible to perform the video dimensioning step and determine if the desired video delivery method is compatible with the other system attributes.

Multicast video streaming requires significant additional configuration, some of which was covered in [Chapter EC-3: Airtime Management](#). Details about other video-specific configuration will be detailed here.

General Multicast Configuration

End-to-end QoS marking described for unicast video is required for multicast video.

[Chapter EC-3: Airtime Management](#) contains a variety of general optimizations for multicast. Video is not the only multicast-based service in WLANs, so it is important to configure these items even if video is not in use:

- Enable IGMP proxy.
- Disable broadcast-filter-all.
- Disable IPv6 to limit unnecessary multicast traffic.
- Configure DSCP-to-WMM mappings.
- Enable dynamic multicast optimization (DMO) with a high threshold of 60 or 80 stations.
- Enable multicast rate optimization (MRO) for nonvideo multicast traffic.

Set a Fixed Multicast Rate for Video Traffic

In addition, as of ArubaOS 6.4.2.3, a feature was introduced to set a fixed multicast rate only for video traffic. This rate is similar to the fixed beacon rate.

This feature is complementary with MRO. MRO applies to nonvideo traffic, but the fixed multicast rate affects only video. The stateful firewall in the controller is used for deep packet inspection to decide which feature to apply on a packet-by-packet basis.

Configure this rate to the highest value that successfully delivers video to most of your users. Aruba recommends starting at 36 Mbps. This rate assumes that the RF design is delivering high enough SINR to all subscribed multicast clients to support the rate. If 36 Mbps goes well, consider trying 48 Mbps. If 36 Mbps does not go well, only then try 24 Mbps. Remember, the lower the video data rate, the less overall capacity the network has for all traffic types.

Table EC6-4 Profiles and Settings for Fixed Video Multicast Rate

Profile	Option	Setting	Reason
SSID Profile	multicast-rate	24 or 36	Set data rate for all multicast traffic.

Use Forward Error Correction

When employing multicast video, it is essential that the video head-end and the mobile player software on the user device support some form of forward error correction (FEC). The exact FEC chosen depends on the manufacturer of the video system.

FEC is necessary because multicast is unacknowledged in 802.11. Therefore, lost frames in a group of pictures (GOP) immediately produce visible artifacts or freeze the video playback until the next key frame is sent. Depending on GOP size configured, this time could be several hundred milliseconds (longer if the wireless network is very congested and multiple key frames are lost).

FEC is applied and removed at the endpoints, and so it is independent of the wired or wireless network.

Video Upload

Considering the limited capacity for downstream video, Aruba strongly recommends that upstream video should not be sent wirelessly if at all possible. (At least not in unlicensed 5-GHz or 2.4-GHz bands.)

Video cameras in VHD areas filming events should be wired. Unlike the mobile clients that can get a low bitrate, these video sources are filmed at 1080P, 4K, or higher resolution with high frame rates. Therefore, they can impose dramatically higher bandwidth requirements on the network and disproportionately affect capacity for all Wi-Fi users.

If remote cameras absolutely have to be wireless due to a mobility requirement or inability to extend a cable, then you have to reserve one or more house channels, which reduces capacity for everyone else. Work with the equipment operator to ensure that it is properly configured to actually use those channels.

Aruba recommends you pursue non-Wi-Fi based wireless technology in a different frequency band for these applications.

Further Assumptions

In this chapter, we have not considered two-way videoconferencing systems. These systems are not common in most VHD areas, with large airports or passenger terminals being the obvious exception.

All of the points about video dimensioning and optimal configuration also apply to two-way videoconferencing systems.

Though it has not been stated, we further assume that video users are mostly stationary in their seats. Generally we do not worry about roaming video in VHD areas. Most customers primarily care about video working once the user population is stationary. It is challenging enough to deliver high-quality video to large numbers of devices, and to add the overhead and latencies associated with roaming is unrealistic. As the wireless architect, you must help educate your customer about what is and is not reasonable to include in acceptance testing procedures in this regard.

That said, Aruba has found that performing video roaming tests can be an excellent diagnostic and tuning tool on the general Wi-Fi system. It is more convenient to watch a video stream on a cart full of tablets than to hold a device continuously to one's ear listening to an audio stream. Video streaming in an **empty** VHD facility is a great tool to rapidly identify optimal AP power levels, evaluate handoff performance, and generally tune the system. However, Aruba does not recommend that video roaming performance SLAs be established for actual events with the VHD facility full.

Chapter EC-7: Configuration Summary

This chapter collects all of the individual configuration best practices from this guide into a single checklist. The best practices are grouped by profile for convenience.



It is assumed that the reader is familiar with Aruba profile-based configuration and the profile hierarchy. This summary **does not** explain the entire profile sequence that is necessary to enable a specific command. For example, defining a custom HT-SSID profile requires a corresponding entry in the parent SSID profile. If you are not sure what this means, do not attempt the configuration. Consult your local Aruba systems engineer or authorized partner for assistance.

Virtual AP Profile

Config Category	Option	Setting	Reason
Filter broadcast and multicast traffic on air	broadcast-filter	All	Prevent all broadcast and multicast traffic from being transmitted on the VAP. (AirGroup traffic is excepted on a station-by-station basis.)
Reduce chatty protocols; Profile peer-to-peer security	deny-inter-user-traffic	Enabled	Prevent peer-to-peer communication.
Enable 802.11k	dot11k-profile	<Profile Name>	Dot11k profile must be defined in the VAP profile.
Multicast speed enhancement	dynamic-mcast-optimization	Enabled	Convert multicast traffic to unicast below the optimization threshold.
	dynamic-mcast-optimization-threshold	80	Specify multicast subscriber count cutoff above which transmission reverts to multicast.

SSID Profile

Config Category	Option	Setting	Reason
Remove low data rates	a-tx-rates	18 24 36 48 54	Eliminate the 6 Mbps and 12 Mbps rates. Consider eliminating the 18 Mbps rate as well if your RF design will support it. Your lowest TX rate should be one lower than your beacon rate.
	g-tx-rates	18 24 36 48 54	Same as a-tx-rates.
Increase base rate of control and management frames	a-basic-rates	24 36	Set minimum control and management frame rate on 5-GHz band to 24 Mbps.
	g-basic-rates	24 36	Set minimum control and management frame rate on 2.4-GHz band to 24 Mbps. Eliminate 802.11b rates.
	a-beacon-rate	24 or 36	Set beacon data rate on 5-GHz band.
	g-beacon-rate	24 or 36	Set beacon data rate on 2.4-GHz band.
Associated device capacity	max-clients	250	Default value of 64 is too small for VHD areas.
Multicast speed enhancement	mcast-rate-opt	Enabled	Activate dynamic multicast rate selection based on the highest “common” data rate.
	multicast-rate	24 or 36	Set data rate for all multicast traffic.
QoS	wmm	Enabled	Enable Wi-Fi multimedia. (enabled by default in 802.11ac.)
	wmm-vo-dscp	56	Set explicit DSCP-to-WMM queue mappings (These are not set by default.)
	wmm-vi-dscp	40	
	wmm-be-dscp	24	
	wmm-bk-dscp	8	
Reduce unnecessary management frames	local-probe-req-thresh	10	Reduce probe response traffic.
Good client distribution	qbss-load-enable	Enabled	Enables the AP to advertise the QBSS load element, which includes: ● Station count: The total number of stations associated to the QBSS. ● Channel utilization: The percentage of time the channel is sensed to be busy. ● Available capacity: The remaining amount of medium time (measured as number of 32us/s) available for a station. A QBSS-enabled device uses these parameters to choose the best AP.
Post-Deployment Site Surveys	advertise-ap-name	Enable	Causes the APs to add an information element that includes the name of each AP in its beacons. Should be disabled when surveys are completed.

HT-SSID Profile

Config Category	Option	Setting	Reason
Channel bonding	40-MHz-enabled	Disabled	Disable 40-MHz operation on the SSID.
	80-MHz-enabled	Disabled	Disable 80-MHz operation on the SSID.
Increase A-MSDU size to 3	max-tx-a-masdu-count-be	3	Allow 3 MSDUs per A-MSDU in [BE] queue.
	max-tx-a-masdu-count-bk	3	Allow 3 MSDUs per A-MSDU in [BK] queue.
	max-tx-a-masdu-count-vi	3	Allow 3 MSDUs per A-MSDU in [VI] queue.
Remove low data rates	supported-mcs-set	3-7,11-15,19-23	Eliminates MCS0 – MCS2 from the HT rate set

ARM Profile

Config Category	Option	Setting	Reason
Channel planning	Assignment	Maintain	For static channel plans. (Use assignment single-band during initial provisioning to create first channel plan, then switch to maintain to lock it in.)
Channel bonding	40MHz-allowed-bands	None	Disable 40-MHz channel assignment. (Must be done on ARM profile for both radios.)
	80MHz-support	None	Disable 80-MHz channel assignment. (Must be done on ARM profile for both radios.)
Power Differentials	Max-TX-power	6 dB or 9 dB higher on the Dot11a ARM profile than on the Dot11g ARM profile	Enforce power differential. (Must be done on ARM profile for both radios.)
	Min-TX-power	6 dB or 9 dB higher on the Dot11a ARM profile than on the Dot11g ARM profile	Enforce power differential. (Must be done on ARM profile for both radios.)

Config Category	Option	Setting	Reason
ClientMatch	cm-sticky-snr	18 dB	Minimum SNR to avoid being steered.
	cm-lb-client-thresh	50	Minimum client count on an AP before steering will occur.
	cm-lb-snr-thresh	20 dB	Min SNR of candidate AP in order to steer a client.
	cm-sticky-min-signal	-75 dBm	Minimum RSSI of candidate AP to steer a client.
	cm-band-g-max-signal	-10 dBm	2.4-GHz clients with strong RSSI should still be steered to 5-GHz.
	cm-steer-timeout	3 seconds	Number of seconds that non-candidate APs should ignore client being steered.
	cm-max-steer-fails	3	Maximum number of steer attempts before client is marked as “unsteerable”.
	cm-unst-ageout	Enable	Enforce ageout of unsteerable client table.
	cm-unst-ageout-intvl	4 hours	Duration that unsteerable client state will be retained.
	cm-stale-age	600 seconds	How long APs hold on to VBR data.
	cm-dot11v	Enable	Helps clients to roam faster, using 802.11v BSS transition commands instead of deauthentication frames.
Temporary Predeployment ARM Training for Channel Plan (revert these settings at the end of the training period)	client-aware	Disable	Allow AP to change channel even if clients are associated to AP and passing traffic.
	video-aware-scan	Disable	Allow AP to go off-channel to scan even if active video clients are present.
	voip-aware-scan	Disable	Allow AP to go off-channel to scan even if active voice users are present.

Dot11a Radio Profile

Config Category	Option	Setting	Reason
Block low-SINR co-channel interference	cell-size-reduction	6 – 10	Increase CCA idle frequency by filtering low-SINR PLCP preambles and data payloads.

Dot11g Radio Profile

Config Category	Option	Setting	Reason
Enable VHT rates in 2.4-GHz	very-high-throughput-rates-enable	Enabled	Allow 256-QAM data rates in 2.4-GHz for compatible clients.
Block low-SINR co-channel interference	cell-size-reduction	6 – 10	Increase CCA idle frequency by filtering low-SINR PLCP preambles and data payloads.
Removing low data rate interference	dot11b-protection	Disabled	Disable protection for 802.11b clients.

Regulatory Domain Profile

Config Category	Option	Setting	Reason
Use all 5-GHz channels	Valid-11a-20mhz-channel-pair	All allowed channels except 144	Ensure that all allowed channels in your country / regulatory domain are enabled. In some countries, DFS channels are disabled by default.
Eliminate channel bonding	Valid-11g-40mhz-channel-pair	No	Remove all 40-MHz channel pairs (enabled by default).
	Valid-11a-40mhz-channel-pair	No	Remove all 40-MHz channel pairs (enabled by default).
	Valid-11a-80mhz-channel-pair	No	Remove all 80-MHz channel pairs (enabled by default).

Traffic Management Profile

Config Category	Option	Setting	Reason
Use Airtime Fairness	shaping-policy	preferred-access	ATF rate limits the air and is vital to prevent older, slower clients from starving newer, faster clients.

Dot11k Profile

Config Category	Option	Setting	Reason
Enable 802.11k	dot11k-enable	Enabled	Enable dot11k operation.
	bcn-measurement-mode	Active	Client should use active probing to populate BSS table.
	rrm-ie-profile	<Profile Name>	Specify RRM sub-profile.

Radio Resource Measurement IE Profile

Config Category	Option	Setting	Reason
RRM IE Profile	quiet-ie	No	Do not silence channel for measurement reports. Required for interoperability and because VHD channels should never be silenced.

AAA Authentication Server Group

Config Category	Option	Setting	Reason
Converged Secure 802.1X SSID	set role	condition Class value-of	Enables server-derivation rules.
	set vlan	condition Aruba-User-Role contains <CPPM> role set-value <VLAN>	Use one entry for each role+VLAN combination that could be returned by CPPM.
	load-balance	Enabled	Distribute requests across cluster.



It is assumed the reader is familiar with RADIUS server configuration on the Aruba controller. This is not the complete configuration for WPA2-Enterprise SSID.

VLAN Interface

Option	Setting	Reason
interface VLAN <X> ip igmp proxy	gigabitethernet <slot>/<port>	IGMP proxy is applied to the VLAN interface, which in turn references physical ports.

Hardening Checklist

Aruba recommends the following network hardening best practices:

- If possible, do not configure an L3 interface on wireless user subnets (including secure subnets) unless a captive portal is being used and redirect is required. The controller should be L2 only on all wireless subnets into which users can be placed.
- Do not be the default gateway for any user subnet
- Place the DHCP server on an L3 separated subnet and use a helper
- Do not be the IP helper for any user subnet
- Configure Aruba Validuser ACL to allow specific user subnets, and disallow protected host IPs or IP ranges including default gateways for each VLAN, DNS, DHCP, etc.
- Guest role should explicitly disallow connection to network infrastructure elements via TCP ports 22 and 4343
- Guest role should explicitly disallow Telnet, SSH and other protocols not required for guest services
- Use a dedicated infrastructure subnet to connect all Wi-Fi controllers, APs, and servers.
- Use ClearPass for administrator authentication using RADIUS and/or TACACS. Monitor the logs.
- Use an IDS solution to monitor both infrastructure and user subnets for suspicious activity
- Shut all unused Ethernet interfaces on the controller
- Monitor for rogue and potential rogue devices in AirWave and on the controller
- Enable “enforce-dhcp” in AAA profiles to prevent users from being able to assign static addresses and gain access to disallowed networks or spoof network resources
- Enable ARP spoof prevention on the default gateway for wireless user subnets, and also on the controller if it has an L3 address on any wireless user VLANs

Appendix EC-A: Worldwide 5-GHz Channel Availability as of March 1, 2015

Channel	Frequency	United States & Canada	Brazil	Europe & Turkey	United Kingdom	Russia	Saudi Arabia	South Africa	Israel	China	Japan	Korea	Singapore	Taiwan	Australia	New Zealand
36	5180	Yes	Indoors	Indoors	Indoors/DFS/TPC	Yes	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors	Indoors	Indoors	Indoors	Yes	No	Indoors	Indoors
40	5200	Yes	Indoors	Indoors	Indoors	Yes	Indoors/DFS/TPC	Indoors	Indoors	Indoors	Indoors	Indoors	Yes	No	Indoors	Indoors
44	5220	Yes	Indoors	Indoors	Indoors	Yes	Indoors/DFS/TPC	Indoors	Indoors	Indoors	Indoors	Indoors	Yes	No	Indoors	Indoors
48	5240	Yes	Indoors	Indoors	Indoors	Yes	Indoors/DFS/TPC	Indoors	Indoors	Indoors	Indoors	Indoors	Yes	No	Indoors	Indoors
52	5260	DFS	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	Yes	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	DFS/TPC	DFS/TPC	Yes	Indoors/DFS/TPC	DFS/TPC
56	5280	DFS	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	Yes	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	DFS/TPC	DFS/TPC	Yes	Indoors/DFS/TPC	DFS/TPC
60	5300	DFS	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	Yes	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	DFS/TPC	DFS/TPC	Yes	Indoors/DFS/TPC	DFS/TPC
64	5320	DFS	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	Yes	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	Indoors/DFS/TPC	DFS/TPC	DFS/TPC	Yes	Indoors/DFS/TPC	DFS/TPC
100	5500	DFS	DFS/TPC	DFS/TPC	DFS/TPC	No	DFS/TPC	DFS/TPC	No	No	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC
104	5520	DFS	DFS/TPC	DFS/TPC	DFS/TPC	No	DFS/TPC	DFS/TPC	No	No	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC
108	5540	DFS	DFS/TPC	DFS/TPC	DFS/TPC	No	DFS/TPC	DFS/TPC	No	No	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC
112	5560	DFS	DFS/TPC	DFS/TPC	DFS/TPC	No	DFS/TPC	DFS/TPC	No	No	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC
116	5580	DFS	DFS/TPC	DFS/TPC	DFS/TPC	No	DFS/TPC	DFS/TPC	No	No	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC
120	5600	DFS ¹	DFS/TPC	DFS/TPC	DFS/TPC	No	DFS/TPC	DFS/TPC	No	No	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC	No	DFS/TPC
124	5620	DFS ¹	DFS/TPC	DFS/TPC	DFS/TPC	No	DFS/TPC	DFS/TPC	No	No	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC	No	DFS/TPC
128	5640	DFS ¹	DFS/TPC	DFS/TPC	DFS/TPC	No	DFS/TPC	DFS/TPC	No	No	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC	No	DFS/TPC
132	5660	DFS	DFS/TPC	DFS/TPC	DFS/TPC	Yes	DFS/TPC	DFS/TPC	No	No	DFS/TPC	No	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC
136	5680	DFS	DFS/TPC	DFS/TPC	DFS/TPC	Yes	DFS/TPC	DFS/TPC	No	No	DFS/TPC	No	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC
140	5700	DFS	DFS/TPC	DFS/TPC	DFS/TPC	Yes	DFS/TPC	DFS/TPC	No	No	DFS/TPC	No	DFS/TPC	DFS/TPC	DFS/TPC	DFS/TPC
144	5720															
Do not use for VHD areas until 802.11ac penetration > 50%																
149	5745	Yes	No	Licensed	No	Yes	No	No	No	Yes	No	Yes	Yes	Yes	Yes	Yes
153	5765	Yes	No	Licensed	No	Yes	No	No	No	Yes	No	Yes	Yes	Yes	Yes	Yes
157	5785	Yes	No	Licensed	No	Yes	No	No	No	Yes	No	Yes	Yes	Yes	Yes	Yes
161	5805	Yes	No	Licensed	No	Yes	No	No	No	Yes	No	Yes	Yes	Yes	Yes	Yes
165	5825	Yes	No	Licensed	No	Yes	No	No	No	Yes	No	Yes	Yes	Yes	Yes	Yes
TOTAL NON-DFS		9	9	4	4	16	8	4	4	9	4	8	9	9	9	9
TOTAL DFS		12 ¹ / 15	15	15	15	0	15	15	4	4	15	12	15	11	13	15
TOTAL		21 ¹ / 24	24	19	19	16	23	19	8	13	19	20	24	20	25	24

1. These channels were temporarily disallowed in 2013-2014 in the US. APs released from 2015 on may use these channels if they pass DFS certification.

Appendix EC-B: 802.11ac Data Rate Table

HT MCS Index	VHT MCS Index	Spatial Streams	Modulation	Coding Ratio	20-MHz		40-MHz		80-MHz	
					800ns	400ns	800ns	400ns	800ns	400ns
MCS 0	MCS 0	1	BPSK	1/2	6.5	7.2	13.5	15.0	29.3	32.5
MCS 1	MCS 1	1	QPSK	1/2	13.0	14.4	27.0	30.0	58.5	65.0
MCS 2	MCS 2	1	QPSK	3/4	19.5	21.7	40.5	45.0	87.8	97.5
MCS 3	MCS 3	1	16-QAM	1/2	26.0	28.9	54.0	60.0	117.0	130.0
MCS 4	MCS 4	1	16-QAM	3/4	39.0	43.3	81.0	90.0	175.5	195.0
MCS 5	MCS 5	1	64-QAM	2/3	52.0	57.8	108.0	120.0	234.0	260.0
MCS 6	MCS 6	1	64-QAM	3/4	58.5	65.0	121.5	135.0	263.3	292.5
MCS 7	MCS 7	1	64-QAM	5/6	65.0	72.2	135.0	150.0	292.5	325.0
	MCS 8	1	256-QAM	3/4	78.0	86.7	162.0	180.0	351.0	390.0
	MCS 9	1	256-QAM	5/6	N/A	N/A	180.0	200.0	390.0	433.3
MCS 8	MCS 0	2	BPSK	1/2	13.0	14.4	27.0	30.0	58.5	65.0
MCS 9	MCS 1	2	QPSK	1/2	26.0	28.9	54.0	60.0	117.0	130.0
MCS 10	MCS 2	2	QPSK	3/4	39.0	43.3	81.0	90.0	175.5	195.0
MCS 11	MCS 3	2	16-QAM	1/2	52.0	57.8	108.0	120.0	234.0	260.0
MCS 12	MCS 4	2	16-QAM	3/4	78.0	86.7	162.0	180.0	351.0	390.0
MCS 13	MCS 5	2	64-QAM	2/3	104.0	115.6	216.0	240.0	468.0	520.0
MCS 14	MCS 6	2	64-QAM	3/4	117.0	130.0	243.0	270.0	526.5	585.0
MCS 15	MCS 7	2	64-QAM	5/6	130.0	144.4	270.0	300.0	585.0	650.0
	MCS 8	2	256-QAM	3/4	156.0	173.3	324.0	360.0	702.0	780.0
	MCS 9	2	256-QAM	5/6	N/A	N/A	360.0	400.0	780.0	866.7
MCS 16	MCS 0	3	BPSK	1/2	19.5	21.7	40.5	45.0	87.8	97.5
MCS 17	MCS 1	3	QPSK	1/2	39.0	43.3	81.0	90.0	175.5	195.0
MCS 18	MCS 2	3	QPSK	3/4	58.5	65.0	121.5	135.0	263.3	292.5
MCS 19	MCS 3	3	16-QAM	1/2	78.0	86.7	162.0	180.0	351.0	390.0
MCS 20	MCS 4	3	16-QAM	3/4	117.0	130.0	243.0	270.0	526.5	585.0
MCS 21	MCS 5	3	64-QAM	2/3	156.0	173.3	324.0	360.0	702.0	780.0
MCS 22	MCS 6	3	64-QAM	3/4	175.5	195.0	364.5	405.0	N/A	N/A
MCS 23	MCS 7	3	64-QAM	5/6	195.0	216.7	405.0	450.0	877.5	975.0
	MCS 8	3	256-QAM	3/4	234.0	260.0	486.0	540.0	1053.0	1170.0
	MCS 9	3	256-QAM	5/6	260.0	288.9	540.0	600.0	1170.0	1300.0
MCS 24	MCS 0	4	BPSK	1/2	26.0	28.9	54.0	60.0	117.0	130.0
MCS 25	MCS 1	4	QPSK	1/2	52.0	57.8	108.0	120.0	234.0	260.0
MCS 26	MCS 2	4	QPSK	3/4	78.0	86.7	162.0	180.0	351.0	390.0
MCS 27	MCS 3	4	16-QAM	1/2	104.0	115.6	216.0	240.0	468.0	520.0
MCS 28	MCS 4	4	16-QAM	3/4	156.0	173.3	324.0	360.0	702.0	780.0
MCS 29	MCS 5	4	64-QAM	2/3	208.0	231.1	432.0	480.0	936.0	1040.0
MCS 30	MCS 6	4	64-QAM	3/4	234.0	260.0	486.0	540.0	1053.0	1170.0
MCS 31	MCS 7	4	64-QAM	5/6	260.0	288.9	540.0	600.0	1170.0	1300.0
	MCS 8	4	256-QAM	3/4	312.0	346.7	648.0	720.0	1404.0	1560.0
	MCS 9	4	256-QAM	5/6	N/A	N/A	720.0	800.0	1560.0	1733.3

Appendix EC-C: DFS Surveys and Operating Rules

The majority of channels available for Wi-Fi on the 5-GHz band in most countries are governed by dynamic frequency selection (DFS) regulations. DFS channels are a vital weapon in the wireless architect's arsenal when planning a very high-density (VHD) WLAN. Depending on the specific country, up to 16 additional channels are currently available for use with Wi-Fi subject to DFS rules.

In [Chapter EC-3: Airtime Management](#), we strongly advocated using these channels in most countries to provide additional user capacity. Therefore, wireless engineers must understand how to assess feasibility of using DFS channels. You must understand how the rules for DFS channels actually operate because there are important differences from non-DFS channels.

How to Conduct a DFS Survey

A DFS survey is the first step to using these channels. A DFS survey requires an Aruba controller and one or more APs of the same type that will be deployed. The process is relatively simple to perform, and has these steps:

1. Install the controller and create a test Service Set Identifier (SSID) and access point (AP) group.
2. Provision all of the test APs into the AP group. Having more APs greatly speeds up the DFS survey process.
3. Perform AP-specific configuration to statically assign a different DFS channel to each test AP.
 - a. Start with the lowest DFS channel in the country where the survey is taking place (for example, 52 in the United States).
 - b. Use only 20-MHz channel widths for VHD deployments.
4. Leave the APs to dwell for a minimum of 4-6 hours on each channel.
5. If a radar event has occurred, it can be noted from the system log, and you will notice that the AP will be on another channel. Radar events are rare except in very specific types of locations.
6. Increment the channel to the next DFS channel and wait for another 4-6 hours.
7. Repeat steps 5 and 6 until all DFS channels have been scanned.

Repeat the entire procedure in at least two locations in a VHD facility: once on a higher floor and once on a lower floor. Testing in different locations is especially important for outdoor VHD areas like stadiums or amphitheaters. Radar can be seen at a higher elevation that is not visible closer to ground level.

Certain PC-based portable spectrum analyzers on the market as of this writing also claim to have DFS survey capability. However, Aruba strongly recommends that you perform the survey with the exact model of AP you plan to deploy because the DFS detection circuitry improves with each new generation of chipset. Using any equipment other than an Aruba AP may produce a different result than you will ultimately see on the system, particularly with regards to any false positives that may be reported.

Behavior of 5-GHz Client Devices in Presence of Radar

If you do plan to use DFS channels, review how DFS works and what you can expect when radar events occur (Figure EC-C1).




Figure EC-C1 DFS Operation Example

The rules for DFS are slightly different for different countries or regions. The definitions of what kind of waveforms should be considered radar signals and how each signal is classified also vary. But, at a high level, Figure EC-C1 shows the basic rules over a period of time for an AP on Channel A that is configured to use a channel subject to DFS rules:

1. Before initiating any transmission on a DFS channel, the device (can be AP or client) monitors the channel for the presence of radar signals for the channel availability check (CAC) time. In most cases, the CAC time equals a minimum of 60 seconds, but it is increased to a minimum of 10 minutes for channels in the 5,600-5,650-MHz subband in Europe (channels 120, 124, 128, 116+, 120-, 124+, and 128-).
2. If no radar signals are detected during the CAC time, the device can start using the channel.
3. While using the channel, the device that “owns” the connection (typically the AP) continuously monitors the channel for radar signals (in-service monitoring). If a radar signal is detected, the AP issues commands to all clients to stop transmissions and switch to a different channel (the channel switch [CS] announcement). After radar detection, the AP is required to clear the channel within 10 seconds.
4. After a radar signal is detected, the AP blacklists the channel and selects a different channel. If that channel is also a DFS channel, step 1 is repeated. If a non-DFS channel is selected, this process no longer applies. Any blacklisted channels are considered unavailable for a minimum of 30 minutes (non occupancy period).
5. At the end of the non-occupancy period, the AP can remain on its current channel or switch back to the original channel (after it completes a new CAC).

As you can see, APs on DFS channels take longer to come up and users on DFS channels can experience brief service interruptions from radar events. Radar frequencies do not align with 802.11 channels, so such events (in theory) can impact multiple Wi-Fi channels simultaneously.

DFS Fact vs. Fiction

Do not assume that simply because your VHD facility is near an airport, military base, or a body of water with shipping traffic that DFS channels are not usable. There may be no radar installation at all, or radar may be present only on specific frequencies, which leaves all other frequencies available for use.

You may also find that only certain parts of your facility experience radar events. In an outdoor stadium for example, the upper seating levels may be more likely than lower bowl seating to experience radar from a nearby source. An indoor venue may see radar events outside but not inside due to attenuation of signal through the building walls.

A DFS survey tells you which channels, if any, should be excluded from your channel plan. Exclude only those channels that experience daily, continuous or recurring radar events. Infrequent radar events does not justify ruling out the use of those channels.

After the network has been deployed, future radar events show up in the ARM history and the system log. These should be periodically monitored for changes, and it's a good idea to set up a SYSLOG alert specific to radar messages. This should provide you with peace of mind that if an event does occur, that you will know about it.

However, with a clean DFS survey, the risk of channel interruption is very low. Consider that a few of your APs happen to experience a radar event when the facility is full of users. The radar event affects only one or two channels. By definition, many more APs cover the area because you have stacked channels for capacity. So clients have many other choices of APs. In this case, the good of the many greatly outweighs the good of the few.

It is a worthwhile price to pay to occasionally have a small number of clients disconnected momentarily to reduce the overall client load on the non-DFS channels. When you avoid DFS channels, you ensure that every one of your users gets less airtime, all the time. When you use DFS channels, you can cut the average per-channel load by more half or more (depending on the particular country). This reduced load increases capacity and performance, all the time.