
VALIDATED REFERENCE DESIGN



VERY HIGH-DENSITY 802.11ac NETWORKS

Scenario 1: Adjacent
Large Auditoriums

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

Scenario 1 – Adjacent Large Auditoriums	4
General Description	4
VRD Scenario Description	5
Physical Layout	5
Device Usage and Mix	6
Application Usage and Bandwidth Expectation	7
Co-Channel Interference	7
Dimensioning	8
Step 1 – Choose Key Dimensioning Metrics	8
Steps 2 and 3 – Estimate Associated Device Capacity and Subnet Size	8
Step 4 – Estimate AP Count	9
Step 5 – Dimension the WLAN Controllers	9
Step 6 –Dimension the Edge	10
Step 7 –Dimension the Core	10
Step 8 –Dimension the Servers	10
Step 9 –Dimension the WAN	11
Capacity Planning	11
Estimate Total System Throughput	11
Estimate Per-User Throughput	17
RF Design	18
Key RF Design Metrics	18
Choose a Coverage Strategy	18
Choose an AP Model	20
Plan the Channels	20
Reserving a House Channel	20
Recommended Security Model	20
End-to-End Architecture	21
Bill of Materials	22
Wireless LAN	22
Wired Network	23
Controller Configuration	23

Scenario 1 – Adjacent Large Auditoriums

Large auditoriums are one of the most common very-high-density (VHD) use cases. Sometimes they occur as a standalone space, as found in high schools or a corporate headquarters building. More often they are found in groups such as in a university lecture hall building or in convention center ballrooms. Often, teaching and note taking are involved, so visitors generally have multiple devices. All these devices increase the required capacity that the WLAN must be designed to support. This scenario will teach you how to properly dimension and design VHD networks for auditoriums.

General Description

An auditorium is any space with these general attributes:

- The room contains at least 200 seats, and a maximum of 1,000 seats.
- Ceiling height is 3–10 m (10–33 ft).
- Seats are permanently installed (except in ballroom case).
- Seats are usually (but not always) arranged to face a presentation area in the front of the room. An alternate layout could include seating around many tables.
- Often multiple adjacent auditoriums are on the same floor, with signal bleed between the walls.
- Often multiple adjacent auditoriums are on different floors in the same building, with signal bleed between floors.
- RF spatial reuse of channels is impossible due to the relatively small physical volume of space and high level of multipath signals.

The scenario uses the generic term “auditorium,” but those attributes describe a surprisingly diverse number of facilities:

- Lecture halls
- Corporate auditoriums
- Large cafeterias
- Hotel ballrooms
- Small performance spaces
- Movie theaters
- Places of worship

As a result, this scenario is one of the single most common deployment cases for high density Wi-Fi®. Aruba recommends that all readers start with this scenario because it is the most common. When you understand the auditorium scenario, you can apply the design principles to almost any other type of high-density coverage zone. Conversely, if you are unclear about any aspect of this scenario, then you will struggle with more complex environments.

VRD Scenario Description

To make this chapter as realistic and useful as possible, we use as our example a hypothetical hotel with a main ballroom of 500 seats and four smaller ballrooms on either side of 200 seats each. Aruba expects readers to adapt the ideas and methods from this chapter to their own specific situation.

Physical Layout

The users in an auditorium are evenly distributed across the space, whether in evenly spaced rows or around evenly spaced tables. The user density is an average of one user per 1–2 m² (11–22 ft²) depending on room configuration. As many as 24 access points (APs) can be deployed in a single auditorium, depending on the total associated device capacity (ADC) target and the number of allowed channels in the regulatory domain. Available mounting locations are often less than ideal, and aesthetic and cable routing considerations limit installation choices. [Figure S1-1](#) shows a typical hotel ballroom.



Figure S1-1 *Typical Configuration of a Hotel Ballroom*

Figure S1-2 shows the seating density in our hypothetical hotel when configured with a lecture hall format. The rooms open into a common hallway.

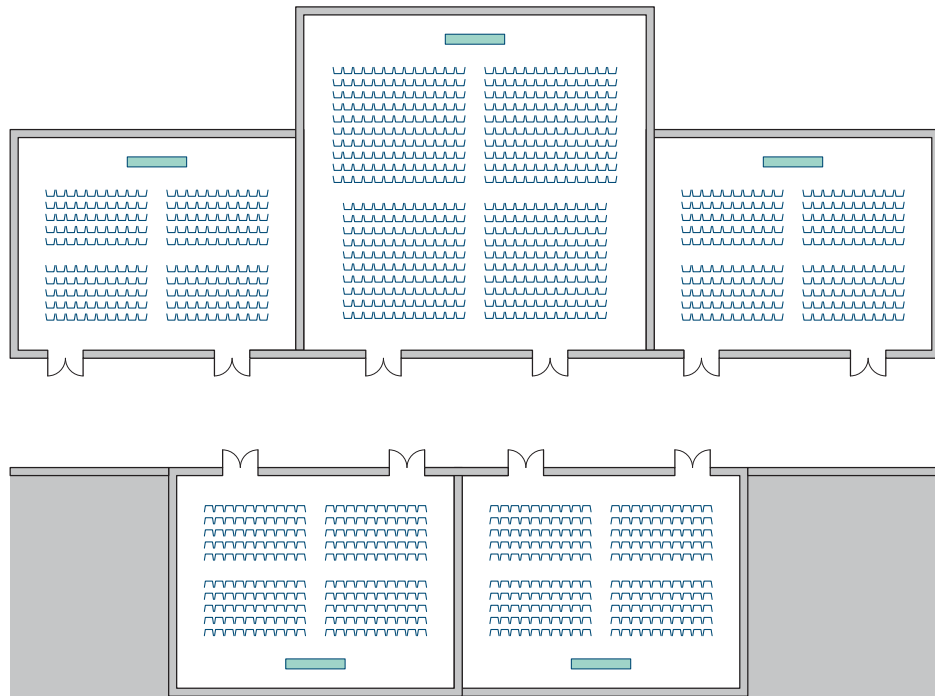


Figure S1-2 Floor Plan for VRD Auditorium Scenario

The user density of a typical auditorium is approximately 20 times greater than an office environment. In a typical office environment with a mix of cubicles and offices, the average density in North America is one person per 10 m² (110 ft²), including common areas. European and Asian countries often have somewhat higher office densities, but they do not approach the density of an auditorium.

Device Usage and Mix

In practice, there are two main profiles of mobile device usage in auditoriums:

- **Work/Study:** The core activity is taking notes, which implies use of a keyboard. Users will have a laptop, netbook, and/or tablet, in addition to a phone.
- **Fan/Guest:** Users are there mostly to watch and listen to some type of performance or presenter. Guests primarily use a smartphone to stay connected to email and social media while in the room. A few users may also have a tablet or a second small Wi-Fi-enabled device.

Using the categories presented in [Table EC2-8 on page 19](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide, [Table S1-1](#) lists the expected spatial stream capabilities of the user devices that apply to the hotel ballroom scenario.

Table S1-1 VHD Spatial Stream Blend Lookup Table

VHD Usage Profile	Devices / Person (Now)	Devices / Person (Future)	1SS (%)	2SS (%)	3SS (%)
Work/Study	3	5	30%	60%	10%
Fan/Guest	1	2	50%	50%	0%

Most auditorium environments have a heterogeneous, uncontrolled mix of client types. The devices are not owned or managed by the facility operator, so they cannot be optimized or guaranteed to have the latest drivers, wireless adapters, or even application versions. Any operating system of any vintage or device form factor could be in use. Wi-Fi radios will be a mix of generations from 802.11a/g through 802.11ac.

Application Usage and Bandwidth Expectation

Auditoriums have a relatively simple set of applications. Users of the wireless network in an auditorium expect moderate throughput, high reliability, and low latency. However, some applications such as email and cloud synchronization may be very high latency without the users being aware of it. The design and configuration of auditorium Wi-Fi coverage must be able to absorb concurrent usage peaks and the initial connection inrush. [Table S1-2](#) lists the network characteristics of common auditorium applications.

Table S1-2 Network Characteristics of Common Auditorium Applications

User Category	Application	Bandwidth	Latency	Duty Cycle
Work/Study	Play courseware (non video)	500 Kbps	Medium	Medium
	Play courseware (video streaming)	1 Mbps+	Low	High
	Test / exam / quiz	Under 250 Kbps	Real-time	Synchronized bursts
Fan/Guest	General internet usage	500 Kbps	Medium	Low
	Email	Under 250 Kbps	High	Low
	Social media	500 Kbps	Medium	Low
	Photo/video cloud sync	1 Mbps+	High	Low

The user traffic in a work/study auditorium WLAN is a variety of application types. The most common applications are web browsing, email, and office collaboration. Except for lecture halls that use e-learning applications, most Wi-Fi users in an auditorium do not expect real-time applications including video or voice communication. Though some users certainly stream multimedia, the vast majority do not.

However, custom applications in the lecture hall scenario can include classroom presentation and exam software, as well as multicast streaming video applications. With the exception of video, these applications are bursty in nature and require concurrent usage by many or all of the wireless clients.

For the fan/guest usage profile, users who record video or audio of sessions in the room do not generally need real-time networking. The media typically is buffered on the device and then uploaded asynchronously to whatever cloud service(s) are configured.

Co-Channel Interference

It is common for several auditoriums of varying capacities to exist side-by-side, or on top of one another on different floors. This situation makes the design aspect more challenging because the rooms are almost always adjacent and close enough to require careful management of co-channel interference (CCI) between auditoriums. This situation can include intended and unintended RF interaction between APs and clients in different rooms.

As a result, such facilities require special RF-design consideration, which is covered in [Chapter P-3: RF Design](#) of the *Very High-Density 802.11ac Networks Planning Guide*. Our hypothetical scenario intentionally considers this situation.

Dimensioning

This section maps the requirements to the nine-step process laid out in [Chapter P-2: System Dimensioning](#) of the *Very High-Density 802.11ac Networks Planning Guide*. Recall that the dimensioning process deals specifically with sizing the entire end-to-end system for a particular number of devices. Dimensioning does not address usable bandwidth.

Step 1 – Choose Key Dimensioning Metrics

[Table S1-3](#) lists suggested values for the key top-level metrics that guide the dimensioning process for the auditorium scenario.

Table S1-3 Top-Level Metrics to Dimension the Auditorium Scenario

Metric	Target
Take Rate	100%
Average devices per person	Work/study – 5 Fan/guest - 2
Associated devices per radio	150
Average single-user goodput	1 Mbps
5 GHz vs. 2.4 GHz split	5 GHz: 75% 2.4 GHz: 25%

In the auditorium scenario, we expect that 100% of users will use the network, and we plan for the worst case that they connect each of their devices. However, later in the chapter we assign different duty cycles to these devices to estimate the potential offered load.

Steps 2 and 3 – Estimate Associated Device Capacity and Subnet Size

Our scenario is a hotel ballroom, so we select the work/study usage profile. We expect each user to have a minimum of three Wi-Fi-enabled devices. However, within the expected lifetime of the network, this number may increase to five devices per person as wearable devices become established. Each user has a primary device that is keyboard-enabled (tablet or laptop), with several secondary devices. [Table S1-4](#) is the planning table for the entire five-room facility.

Table S1-4 ADC and Subnet Plan for Five Ballrooms

Room Number	Seats	ADC (Now)	ADC (Future)	5-GHz ADC (Future)	2.4-GHz ADC (Future)	Minimum Subnet Size
Room A	200	600	1,000	750	250	/22
Room B	200	600	1,000	750	250	/22
Room C	500	1,500	2,500	1,875	625	/20
Room D	200	600	1,000	750	250	/22
Room E	200	600	1,000	750	250	/22
Staff / House	--	25	75	75	0	/24
GUEST ADC	1,300	3,900	6,500	4,875	1,625	/19
STAFF ADC	--	25	75	75	0	/24
TOTAL ADC	1,300	3,925	6,575	4,950	1,625	

The highlighted column in the center is the 5-GHz ADC value, which totals 4,950 devices including both guest and staff devices. This number is based on the 75% / 25% split defined in our key metrics.

Staff devices are estimated at 5 per room today, increasing to 15 per room in the future. This number covers devices carried into the room by staff, as well as permanent “house” devices used for presentation, room control, and other building purposes.

Auditoriums rarely exist alone, and typically all guest users are placed into the same VLAN+subnet wherever they are in the building. So the address space dimensioning that you figure for the guest subnet should include the sum of all the devices you expect in all of the rooms in the facility. Staff devices go into secured VLANs based on user role.

You need not add extra device capacity for the common hallway, as users in the hallway are already counted in the dimensioning plan for controller licensing and IP address space, regardless of which space they are in at any given moment.

Step 4 – Estimate AP Count

To follow the process from [Chapter P-2: System Dimensioning](#) of the *Very High-Density 802.11ac Networks Planning Guide*, we know that AP count is calculated by this formula:

$$AP\ Count = 5\text{-GHz}\ Radio\ Count = \frac{Active\ Device\ Capacity\ (5\ GHz)}{Max\ Associations\ Per\ Radio}$$

We can use this formula to populate a planning table. **Remember that for AP count we must consider the hallway as a separate VHD area.** For planning purposes, we assume that 10% of the total seating capacity might be in the hallway at one time. This increases the 5-GHz ADC from Table S1-4 by 515 devices.

Table S1-5 AP Count for Five Ballrooms

Room Number	5-GHz Guest	5-GHz Staff	Total 5-GHz Devices	Devices per Radio	AP Count
Room A	750	15	765	150	6
Room B	750	15	765	150	6
Room C	1,875	15	1,890	150	13
Room D	750	15	765	150	6
Room E	750	15	765	150	6
Hallway	500	15	515	150	4
TOTAL	5,375	90	5,465		41

Note that the AP count required is higher when calculated on a per-room basis than if you simply divided the total 5,465 5-GHz users by 150 per radio (which yields 37). You must consider each individual VHD area separately for this reason. If you did not do this, some APs would split their capacity across walls.

Step 5 – Dimension the WLAN Controllers

When you dimension controllers, you consider the total device count regardless of how the devices are distributed across frequency bands.

Refer to the in [Table P2-4 on page 20](#) of the *Very High-Density 802.11ac Networks Planning Guide* and use the total future ADC count of 6,575. (The 5-GHz and 2.4-GHz band breakdowns are not relevant to controller dimensioning since each user device consumes a controller license regardless of what channel it uses to connect.) From this value, determine that the minimum controller size for this deployment is an Aruba 7210. That controller is rated for 16,384 users.

If 1+1 high availability (HA) is required, then the customer needs two of these controllers.

Due to the relatively small AP count, it does not make sense to use a dedicated master controller just for the VHD deployment. Aruba recommends a master/local deployment with each controller terminating half of the APs (and serving as failover for the other controller).

If this facility is part of a larger campus with more APs serving more users, then the controllers must be dimensioned to include the entire population of the VHD and non-VHD areas.

Step 6 –Dimension the Edge

To drive the APs in this example, you need 41 full 802.3at gigabit power-over-Ethernet switch ports.

If you prefer to dual-home each AP, you need 82 switch ports. This option requires the AP-225, which has two Ethernet ports. Note that because only 20-MHz channel widths are recommended, it is not possible for the AP to oversubscribe the gigabit Ethernet port. So dual-homing is purely an AP redundancy technique and has no effect on capacity.

Due to the physical layout of the auditoriums, you likely will require two telecom rooms (IDFs) to stay within the 100 m (330 ft) Ethernet distance limit.

Step 7 –Dimension the Core

The ARP cache size and backplane speeds required for this scenario are well inside the scaling limits of most common core switching platforms. Therefore, the details of the core are outside the scope of this auditorium scenario, except to remind you to connect the controllers based on Aruba best practice for a master/local cluster.

Remember that the uplink bandwidth that is required for controller interconnects is double the aggregate system bandwidth. Wi-Fi traffic from the APs is tunneled to the controller before making a hairpin turn back to the edge firewall.

Step 8 –Dimension the Servers

Attendees who visit a lecture hall do not arrive all at one time. Use the arrival profile of a stadium, with peak load of 5% of seating capacity in a five minute interval. For this example, that number is:

$$\text{Peak DHCP Discover Rate} = \frac{6,575 \text{ devices} * 5\%}{300 \text{ Seconds}} = 1.1 \text{ per second}$$

And the DHCP renewal rate is four times this value.

The captive portal peak load should approximate the peak DHCP discover rate, and logins should occur only once per event. The captive portal session time limit should be set so that no one is forced to log in more than once per day.

RADIUS server load is negligible unless you are using 802.11u / PassPoint, which we do not consider in this example.

Step 9 –Dimension the WAN

The WAN uplink bandwidth requirement is equal to the total system throughput (TST), which is calculated in the next section. Note that in many countries, it is easy to exceed 1 Gbps of offered load to the WAN link with a VHD network. We will consider a 21 channel example for the USA that does just that, as well as a 9 channel example that does not.

The TST computed in the example in the next section for a reuse factor of 2.0 is 3,216 Mbps for the 21-channel case. Remember to add a provision for non-Wi-Fi traffic. Therefore, the WAN links should total 3.2 Gbps, which can be achieved with a 10G physical connection that is contractually capped at a lower committed rate.

In the 9-channel case, the TST in the example is 804 Mbps. Therefore, the WAN links should total at least 800 Mbps.



WAN link utilization that is lower than expected can indicate problems with channel plans or simply low usage by the end users. Configure network monitoring platforms such as AirWave to monitor real-time usage as well as record usage over time to identify trends in usage.

Edge firewalls and NAT/PAT devices must be sized for 6,575 individual devices (in addition to any other existing users at this facility in non-VHD areas).

Capacity Planning

To estimate the available system throughput and per-user throughput in this scenario is challenging due to the co-channel interference situation between the adjacent ballrooms. In [Chapter EC-2: Estimating System Throughput](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide, you learned the system bandwidth formula:

$$TST = Channels * Average Channel Throughput * Reuse Factor$$

If we consider each individual room, the largest one requires 13 APs. This number is well below the number of 5-GHz channels available in most countries (see [Appendix EC-A: Worldwide 5-GHz Channel Availability as of March 1, 2015](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide). We are not concerned with a reuse factor inside the ballrooms because the same channel is not used more than once in any single room.

On the other hand, the total number of APs for all five rooms plus the hallway is 41. This number is nearly twice the number of channels that are available in the USA with DFS, and nearly four times the channels that are allowed in many countries if not using DFS. So at a system level, we must be concerned with reuse factor. We have intentionally set up the scenario to expose this issue and how to handle it.

Estimate Total System Throughput

Use the five-step TST process that is defined in [Chapter EC-2: Estimating System Throughput](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide. Begin by establishing the basic facts that are required by the bandwidth estimation process.

Step 1 – Select Channel Count

For completeness, this example considers two channel scenarios.

Table S1-6 Channel Selections for Auditorium Examples

Examples	DFS Usage	5-GHz Channels	2.4-GHz Channels	Total Channels
USA with DFS	Yes	21	3	24
China without DFS	No	9	3	12

Aruba strongly recommends using DFS channels in VHD areas. However, Aruba recognizes that some customers have interoperability concerns, or may be located near radar sources that make using these channels less feasible. So both scenarios are included in this example.



The use of specific country names is simply to make the scenario more realistic. It could be equally valid to say “FCC-based with DFS” and “Non-FCC based without DFS.” Remember that allowed 5-GHz channels are constantly improving for Wi-Fi around the world.

Step 2 – Estimate Unimpaired Multi-client Throughput

In this scenario, we already have defined the usage profile as work/study.

Using [Table S1-7](#) below, choose the blended average bandwidth from the table for the classroom/lecture hall that corresponds to this usage profile. This estimate yields an average channel throughput of 75 Mbps.

Table S1-7 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	60 Mbps	25%	10%	45 Mbps	54 Mbps
Stadium / Arena	40 Mbps	50%	25%	20 Mbps	30 Mbps
Shopping Mall	60 Mbps	10%	5%	54 Mbps	57 Mbps

Step 3 – Apply Impairment Factor

Suggested impairments for both frequency bands and for various VHD facility types are presented in [Table S1-7](#). To keep this example simple, we do not differentiate between frequency bands. Choose a flat 10% impairment.

Applying 10% to the previous step yields an impaired average channel throughput of 67 Mbps. (Use the “2.4-GHz Impaired Goodput” column from [Table S1-7](#)).

Step 4 – Select Reuse Factor

For the moment, consider each room individually. The reuse factor can be equal to 1. We will vary this in a moment when we consider CCI.

Step 5A – Calculate System Throughput Excluding CCI

Table S1-8 calculates the system throughput for both country examples for each individual room and ignores the other rooms.



When you analyze aggregate capacity, do not consider the division between 5 GHz and 2.4 GHz users.

In the USA, even Room C requires fewer channels than the regulatory maximum, so we calculate the full bandwidth. Remember that each room has three 2.4-GHz channels in addition to the 5-GHz channels.

However, in China only 12 channels are available without using DFS. Room C requires one more channel than is available. Therefore, the aggregate bandwidth is capped at the channel limit multiplied by the channel bandwidth. This value is highlighted in **red text** in the table.

Table S1-8 System Throughput Calculation Excluding CCI

Room Number	AP Count	Channels – USA (DFS)	Channels – China (no DFS)	Avg. Channel Bandwidth	Aggregate Bandwidth – USA	Aggregate Bandwidth – China
Room A	6	9	9	67 Mbps	603 Mbps	603 Mbps
Room B	6	9	9	67 Mbps	603 Mbps	603 Mbps
Room C	13	16	12	67 Mbps	1,072 Mbps	804 Mbps
Room D	6	9	9	67 Mbps	603 Mbps	603 Mbps
Room E	6	9	9	67 Mbps	603 Mbps	603 Mbps
Hallway	4	7	7	67 Mbps	469 Mbps	469 Mbps
TOTAL	41				3,953 Mbps	3,685 Mbps

The purpose of looking at each room individually is to help underscore just how significant CCI effects can be between adjacent VHD WLANs. If you have only a single auditorium, co-channel interference is not a big factor and this analysis is all you must do.

Step 5B – Calculate System Throughput Including CCI

In the case of multiple adjacent VHD areas, we cannot ignore co-channel interference. Usually you cannot simply add up the bandwidths of all the individual rooms into a single total value.

For example, here are some of the factors at work in this scenario that would reduce the actual TST below the totals shown above:

- The 2.4-GHz channels are reused in every ballroom. So each room will interfere with all the others most of the time.
- The hallway channels are in use in the ballrooms. So the hallway APs will interfere with the ballroom APs and do not really exist independently.
- All 5-GHz channel numbers are reused at least once. Depending on the free-space distance and the number of walls between same-channel APs, actual achievable reuse will vary greatly.

In this situation, the reuse factor from the TST formula comes in handy. A range of possible system bandwidths can be established by varying the reuse factor value. We know that the worst case scenario is that all rooms interfere with each other. This situation is equal to RF = 1. But depending on the number of channels, APs, and the physical environment, it may be possible to configure the system to minimize CCI and achieve some level of spatial reuse.

Table S1-9 calculates the possible aggregate bandwidths in USA and China regulatory domains for this scenario. From the channel plan in Table S1-9, the maximum number of times that any channel number is reused in 5-GHz is 2.

Table S1-9 System Throughput Calculation Including CCI

Reuse Factor	AP Count	Channels – USA (DFS)	Channels – China (no DFS)	Avg. Channel Bandwidth	Aggregate Bandwidth – USA	Aggregate Bandwidth – China
RF = 1.0	41	24	12	67 Mbps	1,608 Mbps	804 Mbps
RF = 1.5	41	24	12	67 Mbps	2,412 Mbps	No reuse
RF = 2.0	41	24	12	67 Mbps	3,216 Mbps	No reuse

You can see that the reuse factor produces a wide range, particularly with a much smaller channel count. So how do you decide what reuse factor is reasonable?

Unfortunately, in this situation, RF deployment experience is truly required, and simple calculations are not feasible. An experienced WLAN architect can evaluate a fair reuse factor after studying the environment during a site walkthrough. The wall and ceiling construction are the most important factors, followed by the shape and orientation of the individual rooms.

In our VRD scenario, three auditoriums are on one side of a hallway, and two are on the other. The facility layout lends itself to RF spatial reuse if we have enough channels. For simplicity, we will ignore the hallway APs.

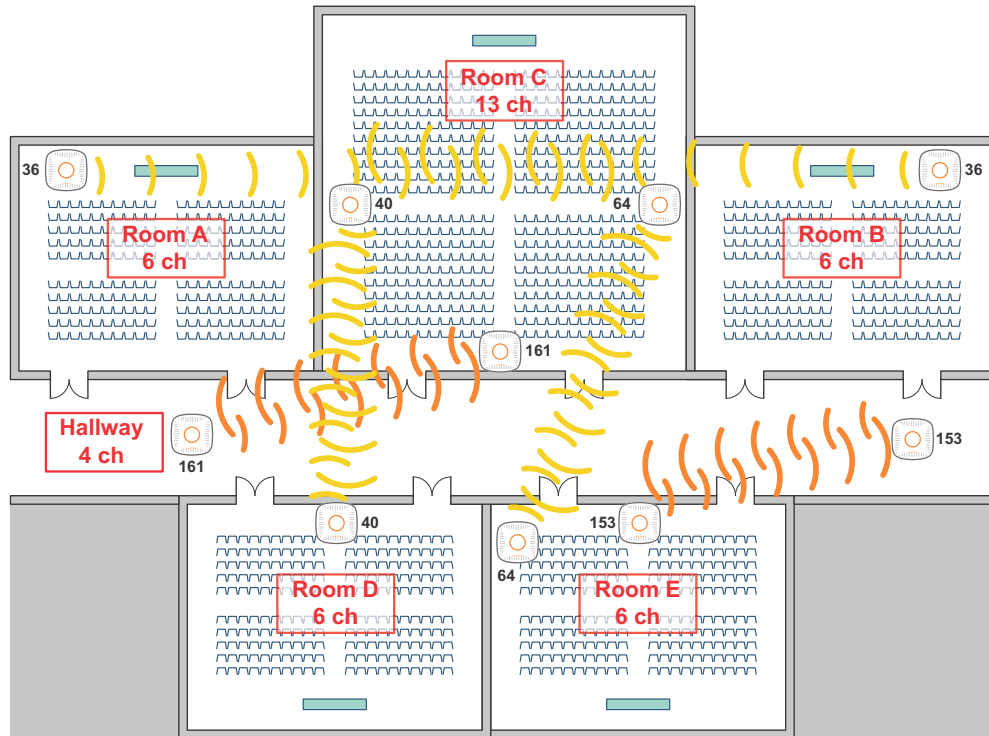


Figure S1-3 Estimated 5-GHz CCI with 21-Channel DFS Plan

Figure S1-3 shows the likely CCI situation in 5 GHz in the USA (or any country with 21 channels), assuming that the channels are laid out according to the following rules:

- No channel is used more than two times
- APs that reuse a channel number must be separated by two walls (except hallway APs)
- APs that reuse a channel number must be physically separated by as much linear distance as possible
- Each room should contribute one channel to the hallway AP pool

Applying these rules produces the channel plan shown in Table S1-10.

Table S1-10 Sample Channel Plan for USA DFS Example

Channel Group	AP Count	2.4-GHz Channels	5-GHz Channels
Room A	6	1, 6, 11	36, 52, 100, 116, 140, 149
Room B	6	1, 6, 11	36, 52, 100, 116, 140, 153
Room C	13	1, 6, 11	40, 48, 104, 108, 112, 132, 44, 56, 60, 64, 161, 165, 136
Room D	6	1, 6, 11	40, 48, 104, 108, 112, 136
Room E	6	1, 6, 11	44, 56, 60, 64, 161, 157
Hallway	4	1, 6, 11	149, 154, 157, 165

Using the terminology from Chapter EC-4: Channel and Power Plans of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide, this plan is a static, local, non-repeating channel plan.

If the channel blocks are assigned as shown in [Figure S1-3](#), no CCI occurs at all between the large center auditorium and its neighbors to the left and right. The auditoriums that reuse channels can be somewhat RF isolated from one another by using the grouping strategy and the building itself. Therefore, in the USA example you could reasonably expect a reuse factor value of 1.5 most of the time. If you utilize Aruba's Cell Size Reduction feature, you could expect reuse of 2.0 most of the time. We only need 41 total APs, and we have 24 channels, so 2.0 is the highest reuse value possible. Our TST equals 3,216 Mbps for the R=2.0 case.

In the 9-channel example, however, things look very different. We do not have enough channels to avoid reuse even inside the largest room. Every neighboring room will interfere with every other room to a significant extent. Therefore, in the 9-channel example, you should not select an RF more than 1.0. This reuse factor gives an aggregate bandwidth of just 804 Mbps!

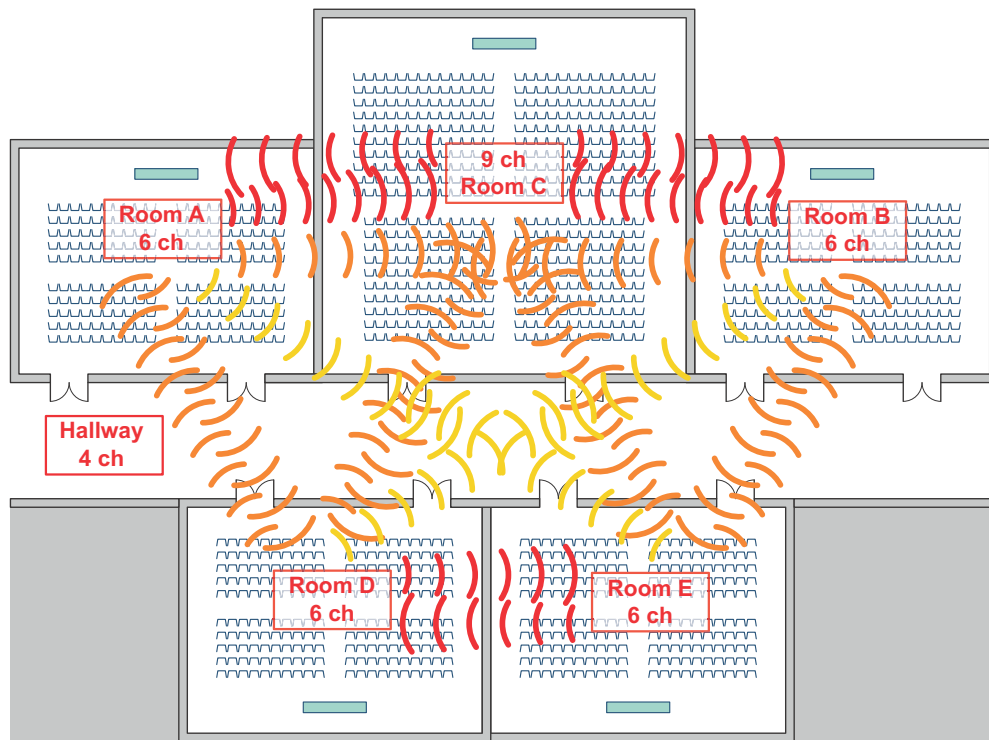


Figure S1-4 Estimated 5-GHz CCI with 9-Channel Non-DFS Plan

The channel plan that is required for a non-DFS deployment of this scenario can be seen in [Table S1-11](#) below. Most channel numbers are repeated five times. The lowest number of reuses is three.

Table S1-11 Sample Channel Plan for China Non-DFS Example

Channel Group	AP Count	2.4-GHz Channels	5-GHz Channels
Room A	6	1, 6, 11	36, 40, 44, 48, 149, 153
Room B	6	1, 6, 11	36, 40, 44, 48, 157, 161
Room C	13	1, 6, 11	36, 40, 44, 48, 149, 153, 157, 161, 165
Room D	6	1, 6, 11	36, 40, 44, 48, 157, 165
Room E	6	1, 6, 11	36, 40, 44, 48, 149, 153
Hallway	4	1, 6, 11	153, 157, 161, 165

This example makes clear one of the most important reasons to utilize extra DFS channels if you have them. Namely, increased physical distance between same-channel APs results in greatly increased RF isolation. Having more channels can produce a non-linear increase in system bandwidth that is significantly greater than you would expect. (However, this effect does not happen in stadiums or arenas where the APs have direct line of sight to each other.)



A USA customer that decides not to use DFS channels has only 9 channels, and therefore is in the second situation.

The preceding analysis of spatial reuse is specific to the 5-GHz band. 2.4-GHz bands have only three nonoverlapping channels, and every individual auditorium and the hallway will use these channels. So in 2.4-GHz, all auditoriums will interfere with each other. Technically, [Table S1-9 on page 14](#) overstates the aggregate bandwidth because the 2.4-GHz channels cannot be reused.

Estimate Per-User Throughput

Now we can estimate average per-user bandwidth. Recall this formula from [Chapter EC-2: Estimating System Throughput](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide:

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

Which can be rewritten as this formula:

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

However, the duty cycle cannot be the same for every device for a user with multiple clients. There is a difference between primary devices that are in a user's hands and secondary devices that are simply in a bag or a pocket. If we treat all devices as primary, then we underestimate the available capacity.

To apply this concept, leverage separate metrics for each type, as shown in [Table S1-12](#).

Table S1-12 Duty Cycle Metrics

Metric	Target
Primary device duty cycle	25%
Secondary device duty cycle	5%

We can easily decide how many primary and secondary devices there are from our future ADC calculation in [Table S1-4](#). In our scenario, each user has five devices by the end of the service lifetime of the network. One of these devices will be considered primary, and four will be secondary. Of the 6,500 total guest devices, this means that 1,300 of them must be primary while 5,200 will be secondary. The 75 staff devices are also assumed to be primary at all times. This brings up our total primary count to 1,375. Now we can plug in our previous work to obtain this formula for the USA example:

$$\text{Average Device Throughput} = \frac{3,216 \text{ Mbps}}{(1,375 \text{ Primary Devices} * \text{Primary Duty Cycle}) + (5,200 \text{ Secondary Devices} * \text{Secondary Duty Cycle})}$$

From the duty cycle value Table EC1-9 in Chapter EC1, select 20% for primary devices and 5% for secondary devices. For the USA, this formula yields this value for the R=2.0 scenario:

$$\text{Average Device Throughput} = \frac{3,216 \text{ Mbps}}{(1,375 * 25\%) + (5,200 * 5\%)} = \frac{3,216 \text{ Mbps}}{344 + 260} = 5.3 \text{ Mbps}$$

For the 9-channel example, we obtain this value:

$$\text{Average Device Throughput} = \frac{804 \text{ Mbps}}{(1,375 * 25\%) + (5,200 * 5\%)} = \frac{804 \text{ Mbps}}{344 + 260} = 1.3 \text{ Mbps}$$

Remember that these numbers are **averages** with comparatively **low duty cycles**. By definition, this analysis assumes that an even distribution of clients attempts to use the Wi-Fi system over time. This assumption is true of many small auditoriums outside the classroom environment. However, the assumption is not valid if you expect any two-way real-time applications. For example, classroom management systems can push content to every primary device at the same time. Online exam systems can push questions and receive answers from all primary devices at the same time. These differences may affect your calculations.

In general, secondary devices can be counted on to behave in average fashion because they are just doing email updates, cloud updates, and so on in the background without user intervention.

RF Design

RF design for auditoriums is very straightforward. You already have the AP counts from the dimensioning process.

Key RF Design Metrics

First, establish the quantitative targets that the RF design must deliver. (See [Table S1-13](#). Remember that these values are net of crowd losses, so when performing acceptance site surveys you would be looking for larger values.)

Table S1-13 Radio Design Metrics

Metric	Target
Minimum cell edge SINR	25 dBm
Minimum cell edge RSSI	-65 dBm

Note that different devices have different radio capabilities, so these values vary significantly across the population. Devices with more than one spatial stream benefit from the 802.11 maximal ratio combining (MRC) feature, which increases SINR. These metrics should be interpreted as applying to the least capable device in the room.

Choose a Coverage Strategy

Overhead or side coverage should be used for auditoriums, as explained in [Chapter P-3: RF Design](#) of the *Very High-Density 802.11ac Networks Planning Guide*. Choose the method that best suits the facility space, aesthetics, and available cabling pathways.

External antennas should never be used in auditoriums of 1,000 seats or less. In general, external antennas do not begin to make a difference until the facility is at least 10,000 seats. These spaces are too small to effectively support spatial reuse and therefore directional antennas are of no value. Antenna suitability should be judged on a case-by-case basis by an experienced RF architect.

Whichever method you choose, APs should be spaced evenly to minimize any adjacent-channel interference. For overhead coverage, that spacing looks like [Figure S1-5](#).

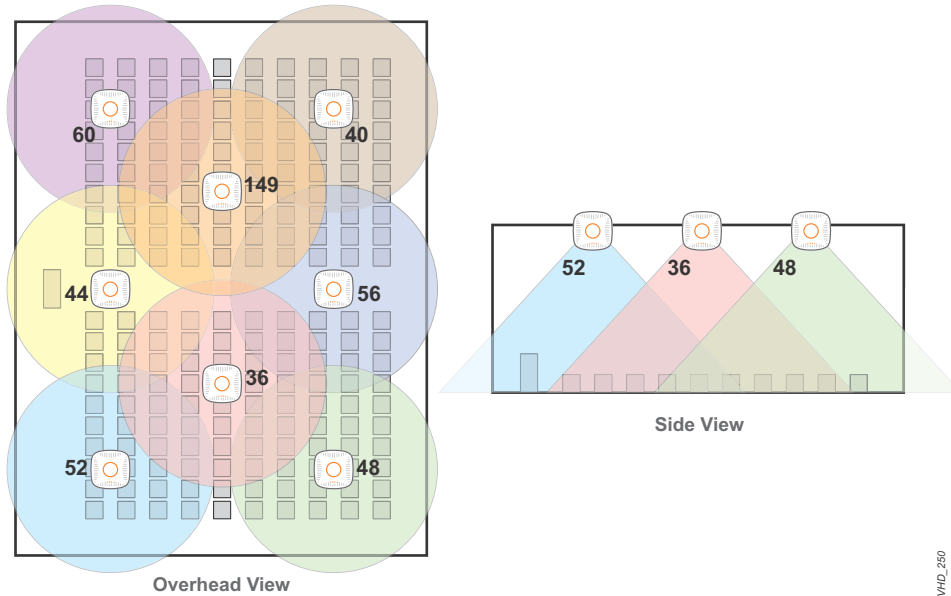


Figure S1-5 AP Placement for Overhead Coverage

Figure S1-6 shows AP placement for side coverage.

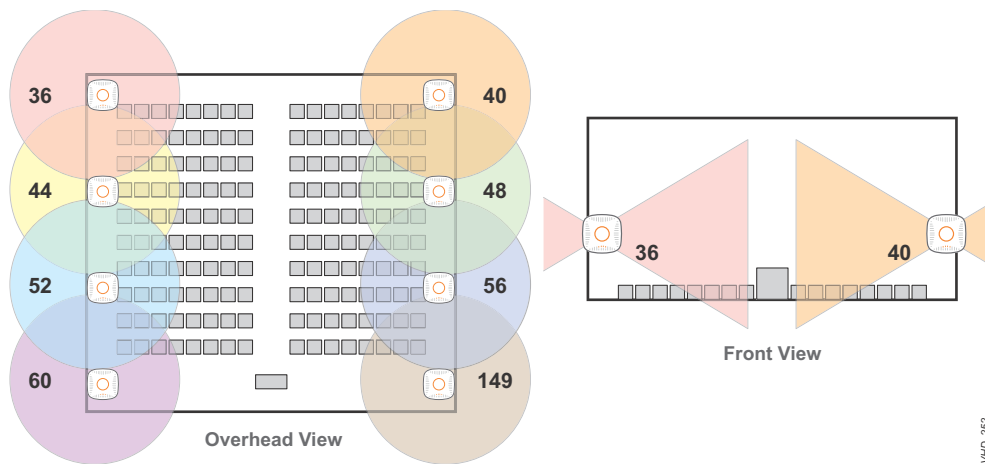


Figure S1-6 AP Placement for Side Coverage

As discussed in [Chapter P-3: RF Design](#) of the *Very High-Density 802.11ac Networks Planning Guide*, be sure to use only one coverage strategy for the entire deployment to ensure even signal dissipation between adjacent rooms. See [Side Coverage with Back-to-Back APs](#) on page 45 in the *Very High-Density 802.11ac Networks Planning Guide* for proper placement guidelines if APs are on both sides of the same wall.

Choose an AP Model

In our scenario, each ballroom has a large number of devices, so Aruba strongly recommends the maximum CPU power possible in each AP. As of this writing, the AP-225 is the highest performing AP that Aruba sells. The AP-225 delivers the best VHD performance.

Plan the Channels

To maximize performance at the room and system level, you must use a static, local, non-repeating channel plan in each room.

- Use each channel only once in each room, if possible (to reduce CCI)
- Adjacent APs should not use adjacent channels, if possible (to reduce ACI)

Non-repeating plans require ARM to be overridden and AP-specific configurations to be applied. Check the regulatory domain profile to enable all of the channels that you plan to use.

Transmit power should be carefully controlled:

- Indoor, low-ceiling auditorium deployments benefit from significant multipath reflection, which require less EIRP.
 - Aruba suggests you try 15 dBm on the 5-GHz band and 9 dBm on the 2.4-GHz band.
 - Increase power if you have clients that experience data rates below the maximum MCS.
- Ensure that 5-GHz is transmitting at least 6 dB higher than 2.4-GHz, whatever your final power levels are. This differential helps clients to self-steer to the correct band.

Reserving a House Channel

If your channel utilizations are low, Aruba recommends that staff and guests share the same channels.

However, if your channel utilizations are high, you may wish to consider holding back one channel for “house” use. This house channel is particularly important if the presenters are using Wi-Fi-enabled equipment and cannot tolerate any interruption in service.

Reserving a channel reduces the overall capacity by taking one full channel out of service, thereby increasing the device load on the other channels. However, if you experience reliability problems for presenters this workaround is useful.

To reserve a channel:

- Remove the desired channel from the regulatory domain profile for the APs serving the room(s).
- Install a dedicated presenter AP near the stage.
- Provision the presenter AP into a separate group, with a regulatory domain profile that has only the reserved channel.
- Configure a secured presenter SSID that is available only on that AP.

In general, reserving a house channel is only possible if DFS channels are being used. This channel can be shared by all VHD rooms in the same facility.

Recommended Security Model

Guest services should be provided via an open SSID that is redirected to a captive portal.

Secure access for facility employees should be provided via an 802.1X SSID using WPA2/AES encryption. Use the facility’s existing AAA server. If the facility is deploying wireless for the first time, ClearPass includes a built-in AAA server that can be used to create secure per-user credentials.

For additional information on the access architecture, see [Chapter EC-5: SSIDs, Authentication, and Security](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide.

End-to-End Architecture

The minimum redundant system-level architecture for this scenario with five ballrooms that serves 6,575 possible devices is shown in [Figure S1-7](#). It is assumed that the core switches in the onsite data center are fully redundant because the service is critical. Having fully redundant switches is convenient because two controllers are required. Per Aruba best practice, controllers should be one-armed to the core with spanning tree disabled. Each controller handles half of the user load when the system is operating normally.

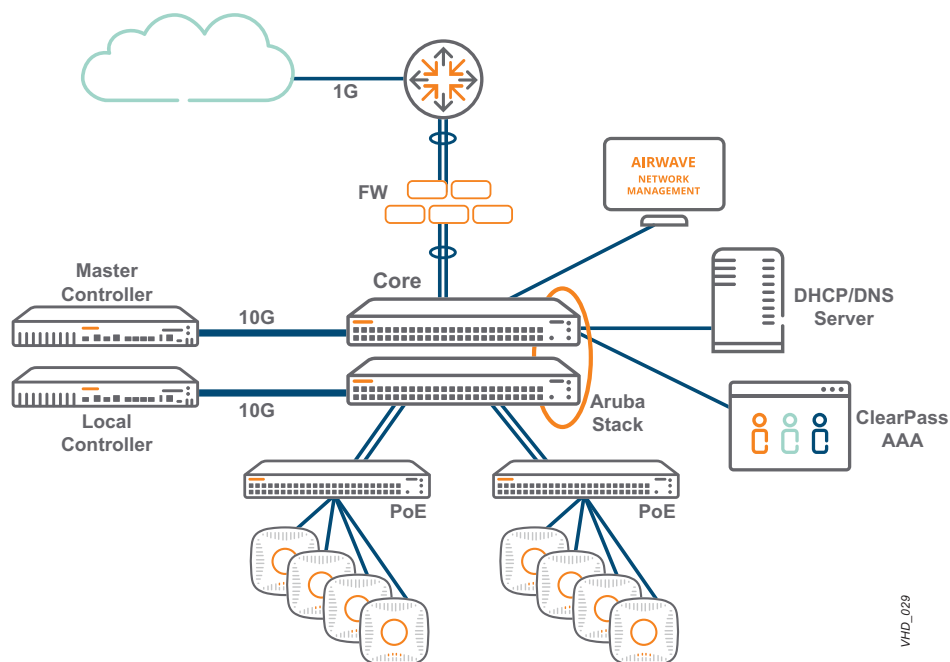


Figure S1-7 Fault-Tolerant VHD System Design

We know from the system bandwidth calculation that we can generate a maximum of 3.2 Gbps of over-the-air goodput for the 24-channel example. Actual load on the controller uplinks will be about 6.5 Gbps for these two reasons:

- Layer 2 tunnel overhead and AP management traffic
- All traffic hairpins through controllers

Therefore, in the USA example, the controllers should use a 10G Ethernet. Aggregating 1G Ethernet links will not be adequate.

In the 9-channel example, maximum system bandwidth is 804 Mbps. Therefore, single 1G Ethernet links may suffice.

In the dimensioning step, we determined that we need two data closets. Each one should dual-home to the core switches.

A high-performance DHCP server with fast transaction speeds is very important. Latency when clients request IP addresses can severely impact performance and always reflects badly on the wireless system.

Captive portal services are provided by a ClearPass Guest server. The captive portal is used to enroll guest users and to obtain their agreement to the facility terms of service. ClearPass must be licensed for 3,900 users immediately (that is, the present capacity, not the future capacity). Additional license capacity can be purchased in the future based on how AirWave reports are trending.

An AirWave server is mandatory because of the multicontroller deployment. AirWave is needed to roll up the user statistics from the separate controllers as well as to provide long term trending. (The controller dashboards show only the last 15 minutes of activity.) Long-term trending is useful to identify typical network usage and spot a slow degradation of service that may be caused by outside interference.

It is a customer decision whether to deploy the DHCP, ClearPass, and AirWave servers in a HA mode.

Moving towards the WAN, the firewall must be sized for NAT/PAT for 6,750 devices with a typical number of sessions.

The minimum link speed from the core all the way to the Internet is 3.2 Gbps in our scenario for the USA DFS example and 800 Mbps for the China non-DFS example.

Bill of Materials

Wireless LAN

Table S1-14 is the BOM for the wireless portion of the network as shown above, and it assumes redundant controllers and nonredundant servers.

Table S1-14 BOM for the Scenario Wi-Fi Network

Part Number	Description	Quantity
7210	Aruba 7210 Mobility Controller with 4x 10GBase-x (SFP/SFP+) and 2x dual media (10/100/1000BASE-T or SFP) ports.	2
PSU-350-AC	350W AC Power Supply. May be used as a redundant power supply or field-replaceable spare for 7200 Series	2
SFP-10GE-SR	10GBASE-SR SFP+; 850nm pluggable 10GbE optic; LC connector; up to 300 meters over multi-mode fiber (Type OM3) or 400 meters with OM4	2
LIC-64-AP	Access Point License (64 Access Point License)	1
LIC-SEC-64	Security Software Bundle (64 AP License) with Policy Enforcement Firewall (PEFNG) and RF Protect (RFP) licenses	1
AP-225	Aruba AP-225 Wireless Access Point, 802.11ac, 3x3:3, dual radio, integrated antennas	41
AP-220-MNT-W2	Aruba Access Point Mount Kit (box style, flat surface wall/ceiling secure mount cradle)	41
AW-50	AirWave license for 50 devices	1
LIC-CP-GM-5K	Guest License for Aruba ClearPass Policy Manager - 5,000 endpoints	1
CP-VA-5K	Aruba ClearPass Policy Manager 5K Virtual Appliance - RADIUS/TACACS+ server with advanced policy control for up to 5,000 unique endpoints.	1

Wired Network

The BOM for the wired core and edge switches shown in [Figure S1-7](#) above is in [Table S1-15](#). It assumes redundant core switches, and edge switches with redundant power supplies. We assume 10G links for the controllers, and aggregated 1G links to each edge switch.

Table S1-15 BOM for the Scenario Wired Network

Part Number	Description	Quantity
S2500-48T	S2500-48T Mobility Access Switch with 48 10/100/1000BASE-T ports plus 4 GbE/10GbE SFP/SFP+ (optics ordered separately)	2
SFP-10GE-SR	10GBASE-SR SFP+; 850nm pluggable 10GbE optic; LC connector; up to 300 meters over multi-mode fiber (Type OM3)	2
S2500-48P	Mobility Access Switch with 24 10/100/1000BASE-T PoE/PoE+ ports and 4 SFP/SFP+ (optics ordered separately)	2
SFP-SX	1000BASE-SX SFP; 850nm pluggable GbE optic; LC connector; up to 270 meters over multi-mode fiber (Type FDDI/OM1)	8

This BOM excludes:

- Server appliances for AirWave and ClearPass
- DHCP/DNS server software and appliance
- WAN edge routers, firewalls, and NAT/PAT equipment

Controller Configuration

The best practice ArubaOS configuration for VHD environments is common to almost all of the VRD scenarios. Read [Chapter EC-3: Airtime Management](#) of the *Very High-Density 802.11ac Networks Engineering and Configuration* guide for detailed instructions.