

MOBILE FIRST CAMPUS FOR MIDSIZE NETWORKS

Design and Deployment Guide

Table of Contents

Introduction.....	1
Purpose of This Guide.....	1
Customer Use Cases	2
Mobile First Campus Design	3
Campus Wireless LAN Design Using Aruba Instant.....	5
Instant Design Components.....	14
Campus Wired LAN Design.....	16
Wired Design Components.....	23
Deploying the Mobile First Campus	26
Campus Wired LAN.....	27
Configuring the Access Switch.....	27
Configuring the Aggregation Switch.....	42
Campus Wireless LAN.....	54
Configuring the Instant Access Point Virtual Controller.....	55
Summary.....	80
Validated Hardware and Software.....	81
What's New in This Version.....	81

Introduction

Wireless has become the primary network access method for today's evolving mobile environments. In the past, wireless networks were a “nice to have,” but they have evolved into a mission-critical lane for connectivity and play a major role in business continuity and in customer and employee satisfaction. In recent years, the number of connected devices per user has increased to more than three, and some estimate it will rise to as many as five per user in the next few years. Employees have their company-supplied PCs, their personal tablets, company-supplied or personal smart phones, and even their smart watches connected to the corporate Wi-Fi network. Users move between locations with their devices and require always-on access. When visiting your employees on-site, guests expect to have access to the Internet from their wireless devices. The Aruba Mobile First Campus network is designed to allow people to move while connected, securely separate employee traffic from guest traffic and to allow enterprises to innovate without being tied to a wired infrastructure. It combines the best wireless products, a wired infrastructure ready to support mobility and Internet of Things (IoT) devices, as well as end-to-end network management with multi-vendor access control.

Because most people work from both company-supplied and personal devices, wireless network access must become ubiquitous to accommodate the new mobile workplace. Guests want Internet access from their personal computers, tablets and smart phones, a desire that becomes a major challenge for IT departments due to the lack of control over the devices. In addition, many IoT devices connect wirelessly to today's networks. As with building control systems, card readers, thermostats, and surveillance cameras, these IoT devices do not have users associated with them. Their traffic is considered machine-to-machine and the devices require machine authentication, which differs from user authentication. Even devices that have traditionally used wired connections, such as shared printers, copy machines, multimedia devices, and high-end workstations, are moving to the wireless world. A network with a few hundred users can easily have over a thousand connected devices.

PURPOSE OF THIS GUIDE

This guide covers the Aruba Mobile First Campus design, including reference designs along with their associated hardware and software components. It contains an explanation of the requirements that shaped the design and the benefits it will provide your organization. The guide describes the access layer as a single system that integrates access points (APs), access switches, aggregation switches, and network management with access-control and traffic-control policies.

Design Goals

The overall goal is to create a simple scalable design that is easy to replicate at different sites in your network. The components are limited to a specific set of products to help with operations and maintenance. The design has a target of sub-second failover when a network device or link between two network devices becomes unavailable. The protocols are tuned for a highly-available network in all functional areas. The design deploys link aggregation and multi-chassis link aggregation between aggregation and access devices. Routed links are utilized at the Core with layer-3 path redundancy.

The guide can be used to design new networks or to optimize and upgrade existing networks. It is not intended as an exhaustive discussion of all options, but rather to present the most commonly recommended designs, features, and hardware.

Audience

This guide is written for IT professionals who need to design an Aruba wired-and-wireless network for a midsize organization with up to 500 users. These IT professionals can fill a variety of roles:

- Systems engineers who need a standard set of procedures for implementing solutions
- Project managers who create statements of work for Aruba implementations
- Aruba partners who sell technology or create implementation documentation

CUSTOMER USE CASES

With so many wireless devices on a network, performance and availability are key. Wireless clients with different capabilities support different performance levels. If the wireless network doesn't self-optimize, slower clients can degrade performance for faster clients. Clients need to intelligently connect to radios on APs to increase network efficiency and performance.

802.11ac Wave 2 Wi-Fi supports speeds greater than 1 Gbps. To accommodate the increased data rates, the APs support 2.5 and 5 Gbps over standards-based, unshielded twisted-pair copper, which works on existing building cabling using Aruba access switches. The access layer acts as a collection point for high-speed wireless devices and must have enough performance to support the bandwidth needs of today as well as scale for the future as the number of connected devices grow

Security is also a critical part of the campus network. Users must be authenticated and given access to the services they need to do their jobs. IoT devices must be identified using machine authentication to prevent rogue devices from using the network. In addition to corporate-managed assets, users connect personal devices, guests need access to the Internet, and contractors need access to the Internet and the organization's internal network.

This type of broad access must be accomplished while maintaining the security and integrity of the network. Connecting so many devices and user types increases the administrative burden, and the network should allow you to automate device onboarding in a secure manner.

Before wireless became the primary network access method, typical network designs provided two or more wired ports per user. It was common to run two network drops to each user's desk and then have additional ports for conference rooms, network printers, and other shared areas, adding up to just over two ports per user. With the trend of users moving to wireless as the primary method of network access, the average wired ports per person is dropping. This trend will continue as more devices move to wireless for connectivity to the network.

This guide will discuss the following use cases:

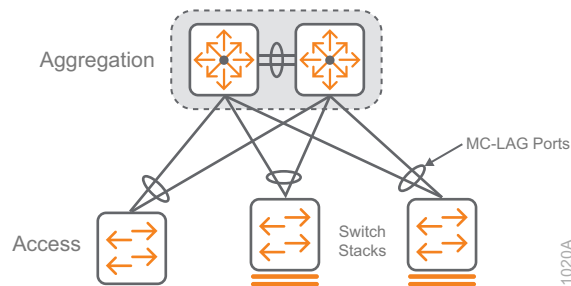
- Wireless as the primary access method for employees
- Wireless guest access for customers, partners, and vendors
- Switch stacking for simplified management, high availability, and scalability
- Link aggregation for high bandwidth, redundancy, and resiliency between switches
- IP multicast to efficiently propagate streaming traffic across the network

Mobile First Campus Design

This design is targeted for midsized organizations supporting up to 500 users with multiple devices per user. The network could be a single building, a few floors in a larger building, or a group of small buildings located near each other. The wireless network requires a common wired local area network (LAN) design which consists of two or three tiers. The access layer is where wired devices and wireless APs connect to the network. The aggregation layer acts as a connection point for multiple access-layer switches. Optionally, the core layer is used to interconnect aggregation-layer switches from multiple buildings or multiple floors in a building.

For a network of up to 500 users, a two-layer campus design is the most common as shown in the following figure.

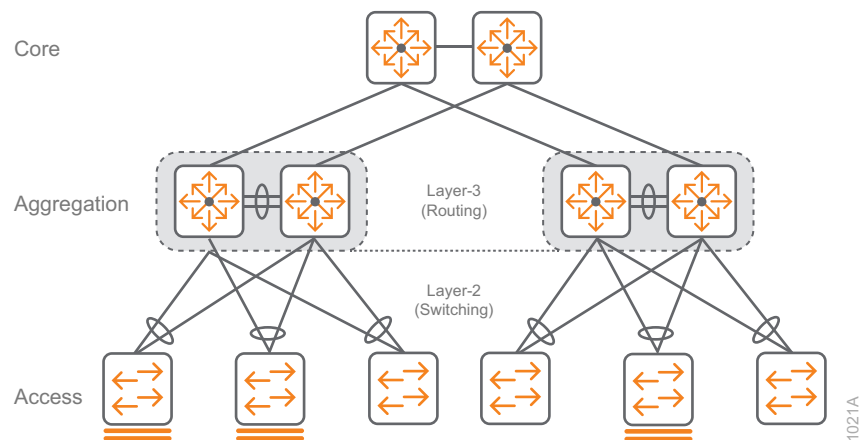
Figure 1 Two-tier campus network



The access-layer switches and switch stacks connect to the dual-switch aggregation using multi-chassis link aggregation (MC-LAG) for higher bandwidth and resiliency.

The three-tier design is used when there are several buildings in a campus that need to be connected and number of aggregation switches or the layout of the physical wiring plant makes more sense to connect everything to a central core. The three-tier campus design is shown in the following figure.

Figure 2 Three-tier campus network



The Mobile First Campus design uses access switches or switch stacks connected to a dual-switch aggregation layer. In networks where 80% or more of the users are connecting via wireless, the number of wired ports in the network is getting close to one per user. This design minimizes the number of different components in order to make operations, maintenance, and troubleshooting simpler.

In this design, Aruba Instant Access APs are used for wireless access because they are easy to deploy and maintain for a network of this size. Both modular and stackable access switches are available, depending on the number of ports needed in the wiring closets. In smaller closets, stackable switches are more cost effective, but at a certain port density, modular access switches will be less expensive than a stack of fixed access switches.

CAMPUS WIRELESS LAN DESIGN USING ARUBA INSTANT

The Aruba Mobile First Campus wireless LAN (WLAN) provides network access for employees, wireless Internet access for guests, and connectivity for IoT devices. Regardless of their location on the network, wireless devices have the same experience when connecting to their services.

The benefits of the mobile-first wireless campus include:

- Location-independent network access improves employee productivity
- Hard-to-wire locations receive network connectivity without costly construction
- Wireless is plug-and-play, with wired LAN switches preconfigured to recognize APs
- Centralized control of wireless environment is easy to manage and operate
- Reliable wireless connectivity, including complete radio frequency (RF) spectrum management, is available with key Aruba management features

Wireless networks today are engineered based on user capacity needs rather than basic wireless coverage. High-speed, high-quality wireless everywhere in the organization is required for today's mobile-first environments. Each client should be able to connect to multiple APs from anywhere in the network. This enables low-latency roaming for real-time applications and allows the network to adapt during routine AP maintenance or an unscheduled outage. A higher density of APs allows the network to support more wireless devices while delivering better connection reliability.

Aruba Instant

For today's wireless networks, there are two main deployment models: one where APs connect to dedicated controllers and one that is controllerless, which is also known as autonomous mode. Aruba Instant is a controllerless wireless architecture that is easy to set up and that supports robust security features. It includes automatic RF management to ensure the best Wi-Fi connection and granular visibility into applications, which helps prioritize business-critical data, limit or block non-business data, and keep malicious actors off your network. A controllerless design is well suited for smaller deployments where tunneling traffic is not needed because of the size of the network and other more advanced controller-based features are not needed. Unlike other autonomous APs, which require a separate management system, an Aruba Instant cluster distributes certain functions across the APs in the cluster and elects a single AP to act as a virtual controller for the remaining management and configuration functions. An Aruba Instant network can be managed with the built-in administrative GUI, with Aruba AirWave, or with Aruba Central, a cloud-based management platform.

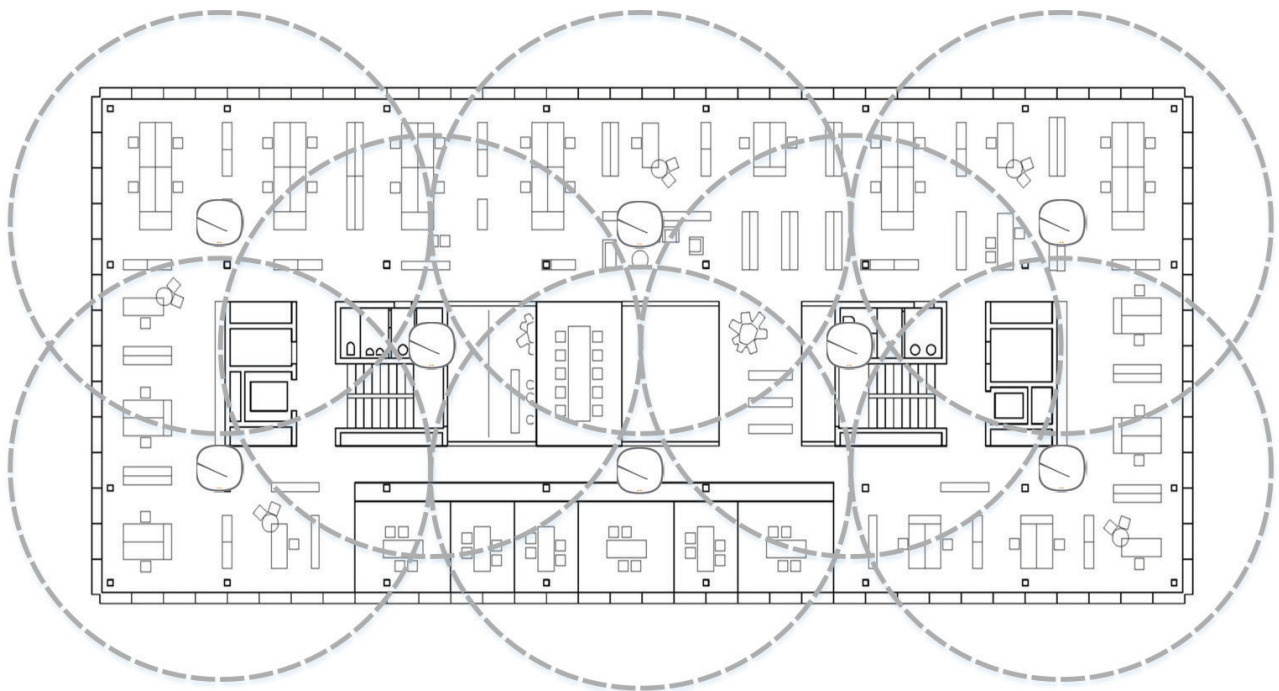
Access Point Placement

For typical capacity in an office environment, we recommend placing APs approximately every 35-50 feet (10-15 meters). Each AP provides coverage for 1500-2500 square feet (140-232 square meters) with enough overlap for seamless client roaming. In traditional offices, the average space per user is approximately 175-200 square feet (16-18.5 square meters), and in open-office environments, the space per user can be as low as 75-100 square feet (7-9.3 square meters). With three devices per user, a traditional office layout with 50-foot AP spacing, and approximately ten users per 2000 square feet, an average of 30 devices are connected to each AP.

The numbers work out roughly the same in higher-density, open-office layouts with 35-foot AP spacing. Because users move around and are not evenly distributed, the higher density allows the network to handle spikes in device count and growth in the number of wireless devices over time. In an average 500-user network with three devices per person, this works out to 1500 total devices, and with 30 devices per AP, this translates to 50 APs. While Aruba Instant can scale well past 50 APs in a single cluster, we are comfortable with an organization building clusters up to 50 APs using standard Instant features. If your design goes past 50 APs per cluster, please involve a skilled wireless partner or Aruba SE/CSE for verification of the solution.

Whenever possible, APs should be placed near users and devices in offices, meeting rooms, and common areas, instead of in hallways. The following figure shows a sample office-floor layout with APs. The staggered spacing between APs is equal in all directions and ensures suitable coverage and seamless roaming.

Figure 3 Sample office AP layout (not to scale)



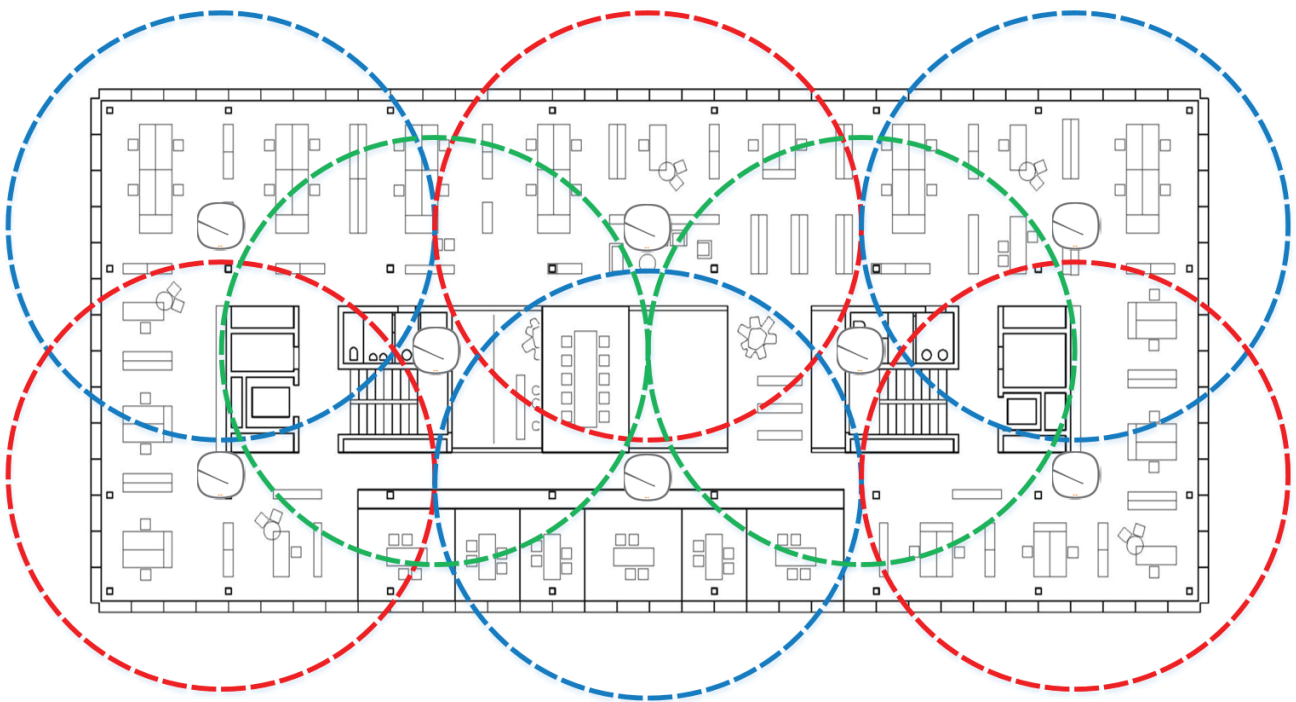
After studying your environment with the 35-50-foot (10-15 meter) rule in mind, make sure you also have enough capacity for the number of users. In an average office environment with APs every 35-50 feet (10-15 meters), the 30 devices per AP average will easily be satisfied. However, if you have high-density areas such as large conference rooms, cafeterias, or auditoriums, additional APs may be needed.

Channel Planning

The Aruba Adaptive Radio Management (ARM) software is very good at automating channel assignment, and for most wireless installations, channel selection and transmit power can be left to its sophisticated algorithms. If you want to plan your channels on your own following the details in this section, please contact an Aruba or partner systems engineer or consulting systems engineer (SE/CSE) for verification of your design.

The following figure shows a typical 2.4-GHz channel layout with each color representing one of the three available non-overlapping channels of 1, 6, and 11 in this band. Reused channels are separated as much as possible, but with only three available channels there will be some co-channel interference which is caused by two radios on the same channel. We recommend only using these three channels for your 2.4-GHz installations to avoid the more serious problem of adjacent channel interference which is caused by radios on overlapping channels or adjacent channels with radios too close together.

Figure 4 Channel layout for 2.4-GHz band with three unique channels



The 5-GHz band offers higher performance and suffers from less external interference than the 2.4-GHz band. It also has many more channels available so it is easier to avoid co-channel interference and adjacent channel interference. Because of the channel advantages, we recommend all capable clients connect on 5 GHz and we recommend converting older clients from 2.4 GHz to 5 GHz when possible. As with the 2.4-GHz spectrum, the radio management software handles the automatic channel selection for the 5-GHz spectrum.

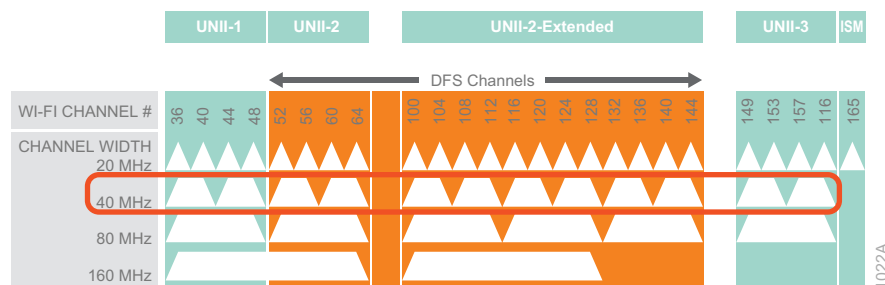
Channel Width

An important decision for 5-GHz deployments is what channel width to use. Wider channel widths mean higher throughput for individual clients but fewer non-overlapping channels, while narrower channel widths results in less available bandwidth per client but more available channels.

In most office environments, 40-MHz-wide channels are recommended because they provide a good balance of performance and available channels. If you are in a high-density open-office environment or you know you will lose channels due to DFS interference, you should consider starting with 20-MHz channels.

However, due to the high number of APs and increasing number of connected devices, there are almost no office environments that would benefit from 80-MHz-wide channels, let alone the much wider 160-MHz channels. The following figure highlights the 40-MHz channel allocation for the 5-GHz band.

Figure 5 802.11ac channel allocation for the 5-Ghz band



Depending on country-specific or region-specific restrictions, some of the UNII-2/UNII-2 Extended Dynamic Frequency Selection (DFS) channels may not be available. In the past, it was common to disable unavailable channels, but today most organizations attempt to use all channels available in their country. In some areas DFS channels overlap with radar systems. If an AP detects radar transmissions on a channel, the AP will stop transmitting on that channel for a time and move to another channel. If specific DFS channels regularly detect radar in your environment, we recommend removing those channels from your valid-channel plan to prevent coverage problems.

Channel 144 is a recent addition to the 802.11ac standard and only newer 802.11ac wireless clients have support for it. If you have older 5-GHz devices, we recommend you disable channel 144.

Using the recommended 40-MHz-wide channels, there are up to 12 channels available. Depending on local regulations and interference from radar or other outside sources, the total number of usable channels will vary from location to location.

A list of the 5-GHz channels available in different countries can be found at the following link: [https://en.wikipedia.org/wiki/List_of_WLAN_channels#5_GHz_\(802.11a/h/j/n/ac/ax\)](https://en.wikipedia.org/wiki/List_of_WLAN_channels#5_GHz_(802.11a/h/j/n/ac/ax))

Spatial Streams

Spatial streaming is a transmission technique in multiple-input and multiple-output (MIMO) wireless communication that allows clients to transmit multiple streams on multiple antennas. The theoretical bandwidth depends on the number of spatial streams and channel width. The following table shows the maximum theoretical bandwidth for the different channel widths and number of available spatial streams.

Table 1 Theoretical bandwidth for 802.11ac at various channels widths and spatial stream counts

Channels Widths	Max Available Channels	1 Spatial Streams (1SS)	2 Spatial Streams (2SS)	3 Spatial Streams (3SS)	4 Spatial Streams (4SS)
20MHz	25	87 Mbps	173 Mbps	289 Mbps	347 Mbps
40MHz	12	200 Mbps	400 Mbps	600 Mbps	800 Mbps
80MHz	6	433 Mbps	867 Mbps	1.3 Gbps	1.73 Gbps
160MHz	2	867 Mbps	1.73 Gbps	2.6 Gbps	3.46 Gbps

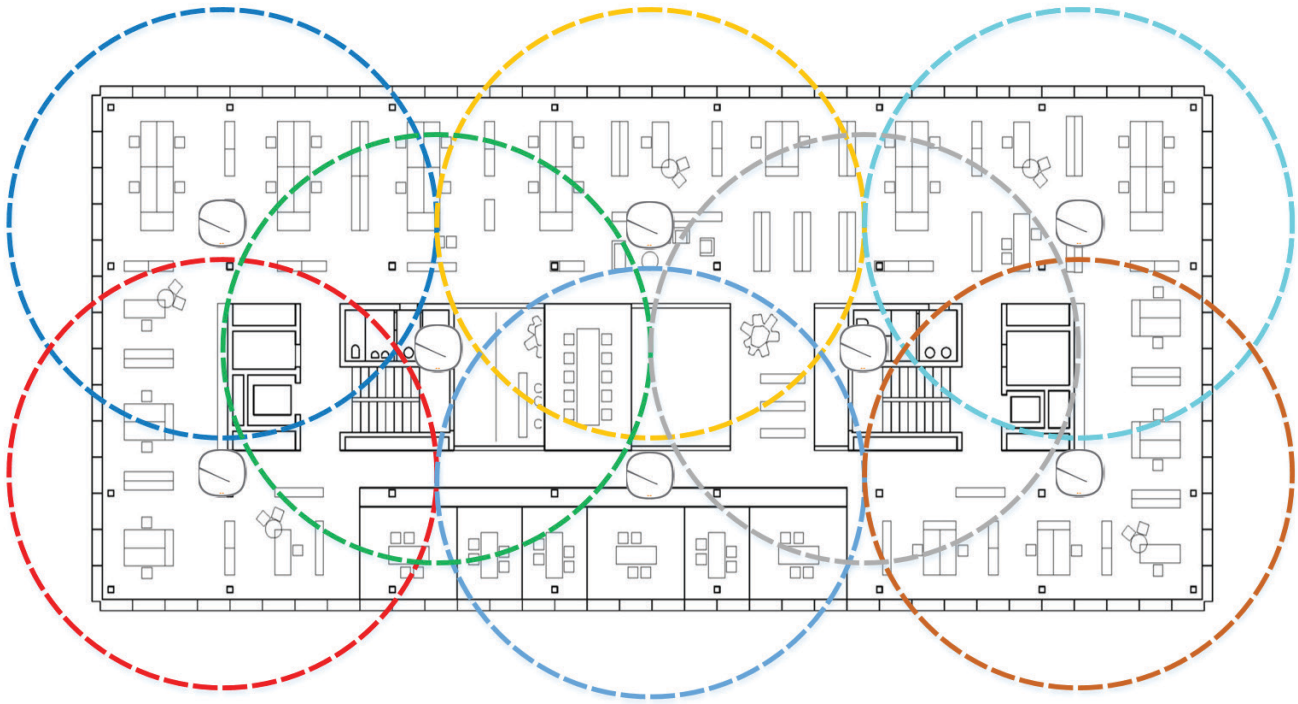
Both the client and the AP need to support the same number of spatial streams to maximize the advantages of this technology. In general, low-power clients like smart phones and low-cost tablets support a lower number of spatial streams and high-power tablets and laptops support a larger number of spatial streams. Aruba ClientMatch balances clients by capability across APs in the network, in order to maximize the service levels available to each type of client.

Channel Summary

The number of APs and their exact placement comes down to performance versus client density. In a high-density deployment, better performance is possible using a larger number of lower-bandwidth channels rather than fewer higher-bandwidth channels. One hundred wireless devices will get better performance split between two radios on 20-MHz channels than they will on one radio using a 40-MHz channel. This is because the more channels you have to use, the better overall throughput will be for a higher number of devices.

A channel layout with eight 40-MHz channels is shown in the following figure. As mentioned previously, a typical Aruba wireless installation will use the ARM software built into the APs for RF channel planning,

Figure 6 Channel layout for 5-GHz band with eight unique 40-MHz channels



Access Point Features

QoS

Quality of service (QoS) allows the network to prioritize traffic so high-priority traffic has preference over low priority traffic while ensuring all applications are treated fairly. With proper QoS, no individual type of traffic can monopolize the network bandwidth. Instead, pre-defined classes ensure that all traffic types are given some amount of bandwidth. Because the access layer is where traffic enters the network from end-user devices, it is important for it to be one of the first policy enforcement points. Traffic entering the network should be classified and tagged based on your organization's requirements.

Management

Aruba ARM and ClientMatch technology are industry leading software features specifically built for enterprise wireless networks. AppRF Technology and Intelligent Application Identification combine to give you unparalleled visibility into the applications running on your wireless network.

ARM—Aruba ARM technology maximizes WLAN performance even in the highest-traffic networks by dynamically and intelligently choosing the best 802.11 channel and transmit power for each Aruba AP in its current RF environment. ARM technology is engineered to address Wi-Fi RF challenges. Leveraging the intelligence embedded in the Aruba infrastructure, ARM has visibility into the entire wireless network and learns about client and application behavior.

ClientMatch—ClientMatch continually monitors a client's RF neighborhood to provide ongoing client band steering and load balancing along with enhanced AP reassignment for roaming mobile clients. This feature is recommended over the legacy band-steering and spectrum load-balancing features, which, unlike ClientMatch, do not trigger AP changes for clients already associated to an AP.

ClientMatch dynamically optimizes Wi-Fi client performance as users roam and RF conditions change. If a device moves out of range of an AP or if RF interference unexpectedly impedes performance, ClientMatch steers it to a better AP. ClientMatch automatically groups multi-user MIMO (MU-MIMO) clients together on Wave 2 APs, so that the AP can transmit simultaneously to multiple clients, thereby realizing the expected Wave 2 throughput and capacity gains for the overall network.

AppRF Technology—Aruba AppRF provides application awareness for thousands of apps, including GoToMeeting, Box, Skype for Business, SharePoint, and Salesforce.com. It also provides web content filtering, enabling IT to control where users can browse on the Internet. The feature uses a cloud database that contains always-up-to-date content and reputation information from millions of web pages.

The AppRF cloud database is updated in real-time with new information about malicious web addresses, enabling AppRF to catch new types of web attacks before they cause damage. To keep clients safe, you can configure them to use the web content filter even when they're not connected to an Aruba Instant network.

Intelligent Application Identification—Aruba's deep packet inspection (DPI) of layer-4 through layer-7 traffic allows the AppRF feature to monitor mobile app usage and performance and to optimize bandwidth, priority, and network paths in real time, even for apps that are encrypted or appear as web traffic. DPI is vital to understanding usage patterns which may require changes to your network design and capacity while identifying many new types of applications:

- Corporate applications such as Box are distinguished from personal applications such as Apple FaceTime, even when they are running on the same mobile device.
- IP multicast video traffic and network services such as Apple AirPrint and AirPlay are automatically prioritized with added policy controls.

- For web-based traffic, DPI resolves the destination address to identify unique applications such as Facebook, Twitter, Box, WebEx, and hundreds of others.
- For encrypted traffic, Aruba AppRF technology uses heuristics to look for traffic patterns and establishes a unique fingerprint to identify those applications.
- For non-business-critical traffic, AppRF can rate limit applications to prevent them from overrunning the remaining bandwidth at a location. Examples would be to limit YouTube and Netflix video streaming for employees during times of congestion. You can also block traffic, such as Netflix video traffic, from your guest network.

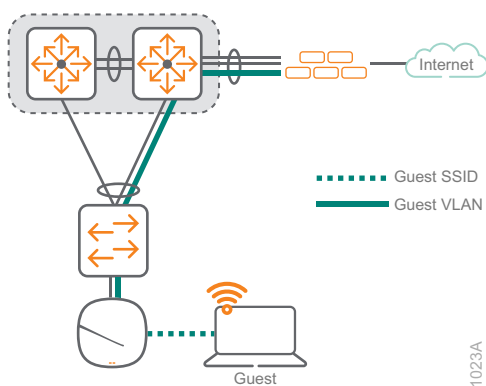
Guest Wireless

Organizations often have a wide range of guests that request network access while they are on-site. Guests can include customers, partners, or vendors, and depending on their purpose, can vary in the type of devices they use and locations they visit in your organization. To accommodate the productivity of this diverse range of guest users and their specific roles, you should deploy guest access throughout the organization and not only in lobby or conference room areas.

The flexibility of the Aruba Mobile First Campus architecture allows the wireless network to provide guest and employee access over the same infrastructure. This integrated ability simplifies network operations and reduces capital and operational costs. The critical part of the architecture is to ensure that guest access does not compromise the security of the corporate network.

Every AP can be provisioned with controlled, open access to wireless connectivity and the Internet. From the wireless AP, guest traffic is placed into a separate VLAN with strict access to the Internet only. For maximum security and for a simplified overall design, traffic is passed from the wireless guest network VLAN to the firewall protecting the organization's private assets, as depicted in the following figure.

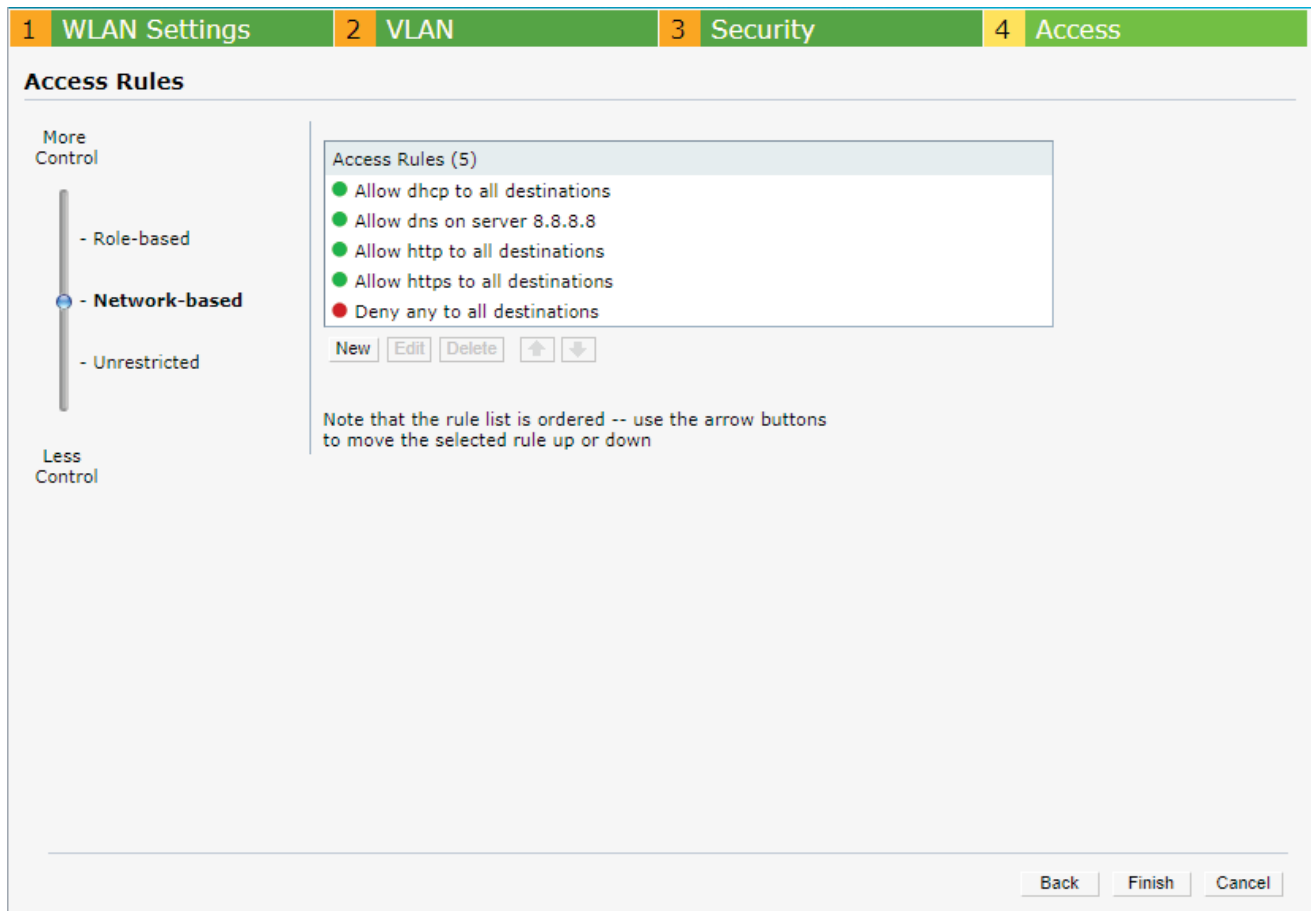
Figure 7 Guest wireless network



1023A

The following example shows how the access rules are applied in the Aruba Instant virtual controller to allow a guest user access to DHCP, DNS, and HTTP/HTTPS network services, while denying the guest user access to all other internal destinations and protocol types.

Figure 8 WLAN guest access rules



To control connectivity, guest users are redirected to a captive portal and must present a username and password to connect to the guest network. The captive portal can be on the Instant virtual controller or an external device. Because the guest traffic must pass through the firewall, strict rules are applied to prevent guest access to the internal corporate network. Lobby ambassadors or other administrative staff can assign temporary guest accounts that require a new password on a timed basis. This design provides the flexibility to tailor control and administration to the organization's requirements while maintaining a secure network infrastructure.

Using the organization's existing WLAN provides a convenient, cost-effective way to offer Internet access for visitors and contractors. The wireless guest network has the following functionality:

- Provides Internet access to guests through an open wireless Service Set Identifier (SSID), with web access control in the firewall
- Segments guest traffic at the AP
- Supports the creation of temporary guest authentication credentials that are managed by an authorized internal user
- Keeps traffic on the guest network separate from the internal network in order to prevent a guest from accessing internal resources

Campus Wireless LAN Design Summary

The Aruba Mobile First Campus WLAN provides network access for employees, guests, and IoT devices. Regardless of their location, wireless devices have the same experience when connecting to their services.

The benefits of the mobile-first wireless campus include:

- Seamless network access for employees, guests and IoT devices
- Plug and play deployment for wireless APs
- Adaptive Radio Management and ClientMatch technology to maximize WLAN performance by dynamically choosing the best Wi-Fi channel and transmit power
- AppRF and Intelligent Application Identification to provide visibility into the applications running on the wireless network

INSTANT DESIGN COMPONENTS

Aruba wireless can be deployed in two main modes, controller-based or controllerless. With Aruba's controllerless model called Instant, there is no central controller and the controller functions are distributed among the APs. Instant is typically used in smaller networks or branch sites and scales up to 128 APs per cluster. In this design, we recommend deploying Aruba Instant with up to 50 APs. If you are planning to install more than 50 Instant APs, please contact an Aruba or partner SE/CSE for verification of your design.

Access Points

The latest generation of Aruba APs, the 3xx models, are dual radio 802.11ac Wave 2 APs and support different throughput and client loads to meet different price points. The last digit in the model number denotes the antenna type. If the number is 4, then the AP has connectors for external antennas. If the number is 5, then the AP has internal antennas. For example, IAP-334 has external antennas and IAP-335 has internal antennas. In most office deployments, internal antenna models are preferred.

The following features are common across the 3xx APs:

- Unified AP for either controller-based (ArubaOS) or controllerless (InstantOS) deployment modes
- Dual Radio 802.11ac AP with MU-MIMO
- Built-in Bluetooth Low-Energy (BLE) radio
- Advanced Cellular Coexistence (ACC) to minimize interference from cellular networks
- QoS for app visibility and control

Access Point Options

Aruba 340 Series Access Points—The Aruba 340 Series is the highest performance AP and supports HPE Smart Rate uplink so it can use the full performance of 3.5 Gbps on two 5-GHz bands or 1.7 Gbps in the 5-GHz band and 800Mbps in the 2.4-GHz band, for a combined bandwidth of 2.5 Gbps. This model is ideal for organizations that require very high density and next generation performance for auditoriums, high-density office environments, or public venues. The Aruba 340 series requires ArubaOS and Aruba InstantOS 8.3 software.

- Dual Radio 4x4 802.11ac AP with MU-MIMO
- Optional dual 5-GHz mode supported, where the 2.4-GHz radio is converted to a second 5-GHz radio
- Antenna polarization diversity for optimized RF performance
- HPE Smart Rate and Gigabit Ethernet uplink ports with Link Aggregation Control Protocol (LACP) support for increased capacity
- Hitless Power over Ethernet (PoE) failover between both Ethernet ports

Aruba 330 Series Access Points—The Aruba 330 Series is a high-performance AP and supports HPE Smart Rate uplink so it can use the full performance of 1.7 Gbps in 5-GHz band and 600Mbps in 2.4-GHz band for a combined bandwidth of 2.3 Gbps. This model is ideal for organizations that require high density and next generation performance for auditoriums, high-density office environments, or public venues.

- Antenna polarization diversity for optimized RF performance
- HPE Smart Rate and Gigabit Ethernet uplink ports with LACP support for increased capacity
- Hitless PoE failover between both Ethernet ports

Aruba 310 Series Access Points—The Aruba 310 Series is a medium-performance AP that supports 1.7 Gbps in the 5GHz band and 300 Mbps in the 2.4-GHz band with a single Gigabit Ethernet uplink. This model is ideal for organization who need to support medium density environments, such as schools, retail branches, hotels, and enterprise offices that don't require multi-gigabit performance.

Aruba 300 Series Access Points—The Aruba 300 Series is an entry-level AP that supports 1.3 Gbps in the 5-GHz band and 300 Mbps in the 2.4-GHz band with a single Gigabit Ethernet uplink. This model is ideal for organizations with medium density environments who want the latest technology but don't need the higher level of performance.

CAMPUS WIRED LAN DESIGN

The campus LAN not only provides wired and wireless connectivity for local users but becomes the core for inter-connecting the WAN, data center, and Internet access, making it a critical part of the network. Campus networks require a high availability design to support the mission-critical applications and real-time multimedia communications that drive the organizational operations.

To accommodate growth in the number of devices, network engineers build wired LANs in layers. A typical wired LAN with up to 500 users will have an access layer and an aggregation layer. With the Aruba Mobile First Campus design, trunks between the layers use multiple links that are actively forwarding traffic for a higher-performance network while reducing the complexity involved in traditional two-layer redundant designs. Breaking the LAN design into layers accomplishes several things that are beneficial to your organization.

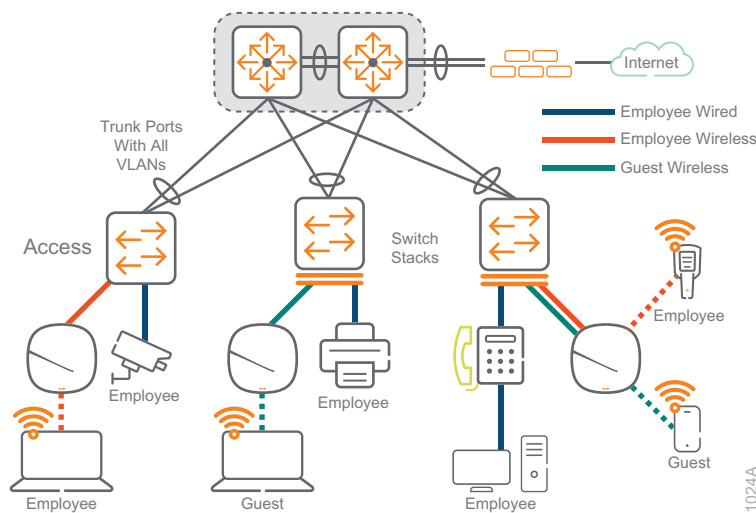
- Limiting functions of the individual layers make the network easier to operate and maintain
- Modular building blocks quickly scale as the network grows
- A repeatable design is faster to deploy across multiple locations

Access Layer

The access layer in this design provides layer-2 connectivity to the network for wired and wireless devices. Because the access layer connects client devices to network services, it plays an important role in protecting users, application resources, and the network itself from human error and malicious attacks. This protection includes verifying the devices are allowed on the network and then making sure the devices cannot provide unauthorized services to end users and cannot take over the role of other devices on the network. The access layer also provides automated services like PoE, QoS, and VLAN assignments in order to reduce operational requirements. The Aruba Mobile First Campus access layer is a layer-2 design to simplify the network as much as possible. Modern LAN switches can easily accommodate thousands of devices, so the old rule of limiting subnets to 254 devices is no longer valid.

Many types of end-user devices connect to the access layer, such as PCs, laptops, smart phones, tablets, and other devices such as printers, video surveillance, and wireless APs. The access layer connects your employee and guest devices into separate VLANs in order to segment the infrastructure into two logical networks. In this design, we use separate VLANs for employee wired, employee wireless, and guest wireless traffic. Employee wired traffic is used by the trusted devices cabled to the LAN switches. Employee wireless is used by trusted devices on Wi-Fi. Employee wired and wireless traffic both have access to all internal resources and the Internet. The guest wireless network is used by untrusted wireless devices, which only have access to the Internet.

Figure 9 Access layer—employee and guest VLANs



Access Layer Switching Features

Stacking

Stacking allows multiple access switches connected to each other through Ethernet connections or dedicated stacking ports to behave like a single switch. Stacking increases the port density by combining multiple physical devices into one virtual fabric, allowing management and configuration from one IP address. This reduces the total number of managed devices while better utilizing the port capacity in an access wiring closet. The members of a stack share the uplink ports, which provides additional bandwidth and redundancy.

QoS

QoS for the wired LAN provides the same benefits for wired clients as was discussed previously for wireless clients on the WLAN. Because the access layer is where traffic enters the network, it is important for it to be one of the QoS first policy enforcement points.

Security Services

Security at the access layer protects end users and the network from configuration errors and malicious attacks. The following security services are recommended at the access layer:

- **Port Security**—Enables you to limit the number of MAC address allowed on a port, stopping MAC flooding attacks. MAC addresses can be learned by the switch or statically configured, and if there is a violation, you have the option of sending an alarm, disabling the port, or both.
- **DHCP Snooping**—Stops IPv4 DHCP starvation attacks, in which an attacker repeatedly requests an address from a DHCP server until no more addresses are available, causing a denial of service to other users. It also prevents rogue DHCP servers by only allowing replies from a trusted server on a trusted switch port, typically the uplink ports to the aggregation layer.
- **ARP Protect**—Stops man-in-the-middle attacks caused by ARP cache poisoning, by verifying the source IP-MAC binding information in the DHCP snooping table. This prevents hosts from sending spoofed ARP messages to fool devices into sending traffic to the wrong address.
- **Dynamic IP Lockdown**—Stops devices from forging their source IP address by inspecting the IP-MAC binding information in the DHCP snooping table. This prevents hosts from injecting traffic into the network to bypass security based on IP source address or to hide their location by forging their source IP address.
- **BPDU Protection**—Prevents loops in the network by putting a non-trunk port into a disabled state for a specified amount of time when it receives a BPDU from another device. This is normally caused by a rogue device being connected to an access port on a switch.

- **DHCPv6 Snooping**—Stops IPv6 DHCP starvation attacks, in which an attacker repeatedly requests address from an IPv6 DHCP server until no more addresses are available, causing a denial of service to other users. It also prevents rogue IPv6 DHCP servers by only allowing replies from a trusted server on a trusted switch port, typically the uplink ports to the aggregation layer.
- **IPv6 RA Guard**—Stops rogue IPv6 clients from advertising themselves as routers. The IPv6 RA Guard feature on the switch analyzes Router Advertisements (RA) and filters out the ones sent by unauthorized devices.

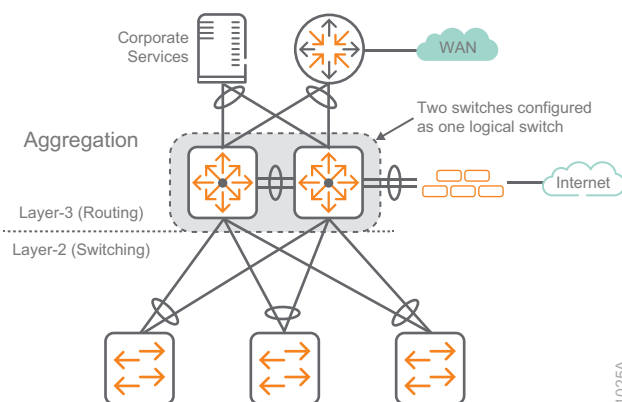
IP Multicast

The access layer switches use a key IP multicast feature called IGMP snooping. IGMP snooping is designed to prevent hosts on a local network from receiving traffic for a multicast group they have not explicitly joined. The feature provides layer-2 switches with a mechanism to prune multicast traffic from ports that do not contain an active multicast listener.

Aggregation Layer

The aggregation layer acts as the boundary between layer-2 switching and layer-3 routing. The aggregation layer provides layer-3 services, routing LAN traffic between networks in the campus and out of the campus to other networks across the WAN. Because layer-2 networks are terminated at the aggregation layer, it segments the network into smaller broadcast domains. As more access layer switches are added, it becomes difficult to interconnect them with a full mesh because meshing uses the uplink ports quickly and daisy-chaining limits the overall performance of the network. The aggregation layer increases network scalability by providing a single place to interconnect the access layer switches, giving you high performance and single hop connectivity between all switches in the aggregation block. The aggregation layer also becomes the ideal location for connecting other network services, such as the WAN aggregation, Internet DMZ, and server rooms for a midsize organization.

Figure 10 Aggregation layer—routing and switching boundary

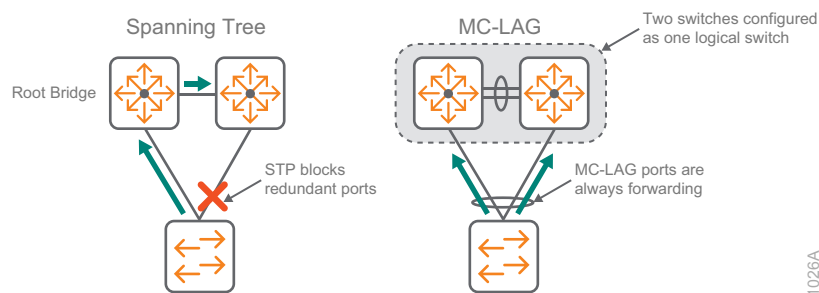


Aggregation Layer Switching Features

Multi-Chassis Link Aggregation

Multi-Chassis Link Aggregation Group (MC-LAG) allows the aggregation layer switch pair to appear as a single device to other devices in the network, such as the access layer switches. MC-LAG allows all uplinks between switch stacks to be active and passing traffic for higher capacity and availability, as shown in the right side of the following figure. Older, redundant designs relied on Spanning Tree Protocol (STP), which blocked redundant links, as shown in the left side of the following figure. It can take up to 50 seconds for a traditional spanning-tree port to transition from blocking to a forwarding state and traffic is not forwarded during the re-convergence time. MC-LAG ports are always forwarding, so the re-convergence time for active traffic on a failed link is less than 100 ms.

Figure 11 Traditional spanning tree vs MC-LAG design



From an STP standpoint, the access to aggregation layer MC-LAG connection looks like a single link, removing all loops in the topology and preventing link or switch failures from causing STP re-convergence.

Depending on the switch model, the Aruba switches support MC-LAG using either backplane stacking or Virtual Switching Framework (VSF), in order to appear as a single switch to other devices in the network.

The benefits of multi-chassis link aggregation are as follows:

- **Performance and Capacity**—A stack creates a pool of network ports with optimized forwarding, so any member of the stack can utilize the shared uplinks in order to meet network demands. MC-LAG combines links from individual switches in the stack, allowing them to act as one connection, which increases the performance of the uplinks.
- **Resiliency and Redundancy**—If a MC-LAG member switch fails, the other member continues to operate, which reduces recovery time. Links to the switches in a MC-LAG group are split across the stack members which provides additional bandwidth, link redundancy, and physical device redundancy.
- **Simplifies Management and Configuration**—Even though a stack consists of multiple physical devices, a stack is managed as a single device with a single configuration, which simplifies network design and management.

IP Routing

In a midsized organization, all departments need to be connected and sharing information. To accomplish this in an easy, scalable manner, a dynamic routing protocol is needed. Open Shortest Path First (OSPF) is a dynamic, link-state, standards-based routing protocol that is commonly deployed in campus networks. OSPF provides fast convergence and excellent scalability, making it a good choice for midsize networks because it can grow with the network without the need for redesign.

OSPF uses areas which provides segmentation of the network to limit routing advertisements and allow for route summarization. Area segmentation is normally done on logical network boundaries, such as buildings or locations, and it helps minimize the impact of routing changes across the network. In large networks with WANs, multiple OSPF areas are very useful, but in campus networks of this size a single area is recommended.

IP Multicast

IP multicast allows a single IP data stream to be replicated by the network and sent from a single source to multiple receivers. IP multicast is much more efficient than sending multiple unicast streams or flooding a broadcast stream that would propagate everywhere. Common examples of multicast traffic in a campus network are IP telephony music on hold (MOH) and IP video broadcast streaming.

This design uses protocol independent multicast (PIM) sparse mode to route multicast traffic on the network. Rather than build a separate routing table, PIM uses the unicast routing table. In our case, the routing table created by OSPF is used for reverse path forwarding. The three mechanisms to route multicast in this design are the rendezvous point (RP), bootstrap router (BSR), and Internet Group Management Protocol (IGMP).

The BSR is elected from a list of candidate-BSRs configured on the network. There can only be a single active BSR on the network. The BSR advertises RP information to all PIM-enabled routers in the network, freeing you from having to statically configure the RP address on each router in the network. The BSR also allows for backup RPs to be configured for multicast groups. If a primary RP fails, the network can switch to the backup automatically. Typically, routers in the core of the network are configured as the BSR candidate routers.

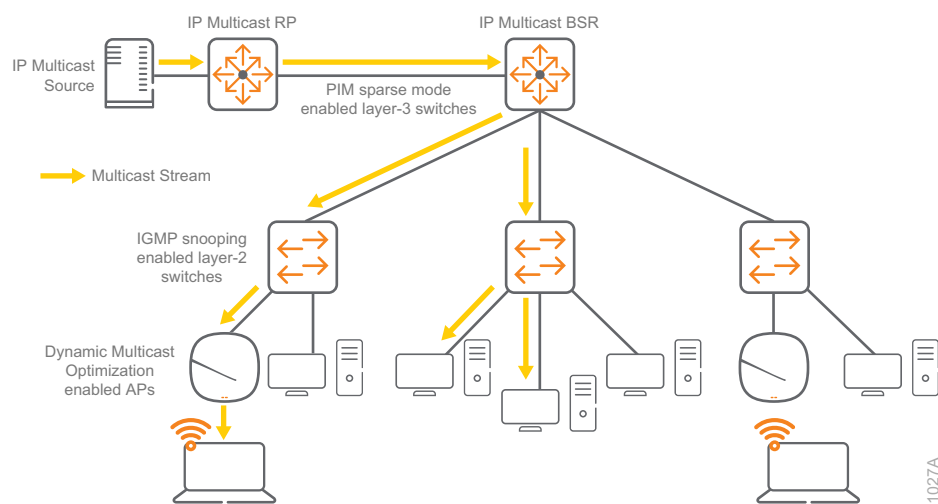
The RP is the root of the multicast tree for multicast traffic using sparse mode. Because it is the root for shared multicast traffic, the RP is normally placed at the core of the network or at the point where the most multicast senders are located. Multiple RPs can be configured for redundancy, although only one RP can be active for a multicast group. Multicast sources are registered to the RP when the local multicast router sends a unicast register message to the RP.

When a client wants to join a multicast group, it sends an IGMP message to its local multicast router. The local multicast router, called the designated router (DR), forwards the join message towards the RP and all routers in the path do the same until the join reaches the RP. Multicast traffic is forwarded back down the shared tree to the client. Periodic join messages are sent to the RP for each multicast group with active clients. If a DR wants to stop traffic from a multicast group because it no longer has active clients it can send a prune message to the RP. To prevent the DR from flooding traffic to all clients on a local subnet, layer-2 switches snoop the IGMP messages and only forward traffic to clients that have sent a join message.

The 802.11 standard states that multicast traffic over WLAN must be transmitted at the lowest basic rate so all clients are able to decode it. We recommend enabling Dynamic Multicast Optimization (DMO) to allow the AP to convert the multicast traffic to unicast for each client device. Unicast packets are transmitted at the higher unicast rate which decreases the airtime utilization and increases overall throughput. IGMP snooping must be enabled on the layer-2 switches for DMO to work.

The following figure shows a multicast source registered with the RP and sending traffic to clients that have joined the multicast group. Note that clients not receiving the multicast stream have not joined the group.

Figure 12 IP multicast with PIM sparse mode, IGMP snooping and DMO



Campus Wired LAN Design Summary

The simplified access and aggregation design provides the following benefits:

- An intelligent access layer provides protection from attacks while maintaining user transparency within their layer-2 VLAN boundaries.
- Redundant uplinks forward traffic, providing higher bandwidth and resiliency without creating layer-2 STP loops in the network.
- The MC-LAG aggregation layer provides reduced complexity while improving recovery times during network failures.
- The aggregation/core layer provides IP routing using OSPF and IP multicast using PIM sparse mode with redundant BSRs and RPs.
- The aggregation/core layer is the logical place to connect critical networking devices such as corporate servers, WAN routers, and Internet-edge firewalls.

WIRED DESIGN COMPONENTS

The wired LAN in the Mobile First Campus uses a hierarchical, modular design. Each layer performs specific functions helping to simplify the design, making the network easier to deploy, manage, and maintain. While there are many hardware choices that will work at the different layers in the network, this design focuses on products that are the most common and easily supported options in each layer of the network, with general guidance on which option to choose.

Access Switches

The access layer connects wired devices to the network, such as APs, workstations, multi-function printers, and other devices that don't support Wi-Fi or that do need higher performance than a wireless connection can provide. The access layer also provides PoE to devices such as APs, IP phones, and IP cameras.

The following features are common across the Aruba access switches:

- Support for security and network management with Aruba ClearPass, Aruba AirWave, and cloud-based Aruba Central
- REST APIs for the software-defined network
- PoE

The number of ports needed in an access closet and the performance required will decide what access switch model is the best fit for your network.

Access Layer Switching Options

Aruba 5400R—The Aruba 5400R chassis supports a variety of interface modules that provide copper and fiber interfaces in different speeds and densities. At the access layer, the switch supports up to 96 HP Smart Rate Multi-Gigabit or 288 1-GbE ports with PoE+. This switch is ideal for organizations that need large numbers of access ports in high density areas of their network (majority of access closets with 96+ ports).

- Layer-3 modular switch with VSF stacking, tunnel node, ACLs, robust QoS, low latency, and resiliency
- HPE Smart Rate for high-speed multi-gigabit bandwidth (IEEE 802.3bz) and PoE+
- Scalable line-rate 40 GbE for wireless traffic aggregation

Aruba 3810M—The Aruba 3810M is available with either 24 or 48 1-GbE access ports with PoE+ (30W) on each port and either 4 SPF+ ports or 2 40-GbE ports on an optional expansion module. The 3810M is also available in a model with 40 1-GbE ports and 8 HPE Smart Rate ports capable of 1, 2.5, 5, or 10 GbE. The 3810M supports backplane stacking with up to 10 switches in a single stack and advanced layer-3 services. This switch is ideal for organizations that have larger access closets requiring larger switch stacks, are deploying or planning on deploying 802.11ac Wave 2 APs, and want a switch with high performance and room for future growth.

- Layer-3 switch with backplane stacking, tunnel node, ACLs, robust QoS, low latency, and resiliency
- HPE Smart Rate for high-speed multi-gigabit bandwidth (IEEE 802.3bz) and PoE+
- Modular line-rate 10-GbE and 40-GbE ports for wireless aggregation

Aruba 2930M—The Aruba 2930M is available with either 24 or 48 1-GbE access ports with PoE+ (30W) on each port and either 4 SPF+ ports or 2 40-GbE ports on an optional expansion module. The 2930M is also available in a model with 40 1-GbE ports and 8 HPE Smart Rate ports capable of 1, 2.5, 5, or 10 GbE. The 2930M supports backplane stacking with up to 10 switches in a single stack and dynamic layer-3 services. This switch is designed for organizations wanting to create a digital workplace optimized for mobile users with an integrated wired and wireless access network.

- Layer-3 switch with backplane stacking, tunnel node, ACLs, and robust QoS
- HPE Smart Rate for high-speed multi-gigabit bandwidth (IEEE 802.3bz) and up to 1440 W PoE+
- Modular 10-GbE or 40-GbE uplinks
- Models with 24 ports of HPE Smart Rate with IEEE 802.3bz

Aruba 2930F—The Aruba 2930F is available with either 24 or 48 1-GbE access ports and 370W PoE+. The switch supports VSF, allowing you to stack up to 4 switches using available front ports. While the 2930F supports basic layer-3 features, it is typically deployed as a layer-2 switch. This switch is ideal for organizations that have smaller access closets requiring only one or two switches, are looking for good performance, and who can accept a limited feature set in return for lower cost.

- Layer-3 switch with VSF stacking, tunnel node, ACLs, and robust QoS
- Convenient built-in 1GbE or 10GbE uplinks and up to 740 W PoE+

Aggregation Switches

The aggregation layer provides connectivity for all the access layer switches and connects to any external networks in the campus LAN. The aggregation layer is responsible for layer-3 routing in this design and it handles all traffic between networks on the campus LAN and traffic leaving the LAN for the data center, the WAN or the Internet. For high availability, the aggregation layer consists of a pair of switches acting as a single switch. If a switch fails or needs to be taken out of service for maintenance, the other switch continues forwarding traffic without interruption to the LAN services.

The following features are common across the aggregation switches:

- HPE Smart Rate for high-speed multi-gigabit bandwidth (IEEE 802.3bz) and PoE+
- Support for security and network management with Aruba ClearPass, Aruba AirWave Network, and cloud-based Aruba Central
- REST APIs for the software-defined network

Aggregation Layer Switching Options

Aruba 5400R—The Aruba 5400R chassis supports a variety of interface modules that provide copper and fiber interfaces in different speeds and densities. The switch supports up to 96 10-GbE ports (SFP+ and 10GBASE-T), 96 HP Smart Rate Multi-Gigabit, or 288 1-GbE ports with PoE+. This switch is ideal for organizations that need to aggregate many access switches and may need to connect servers, firewalls or other network appliances directly to the aggregation layer.

- Layer-3 modular switch with VSF stacking, static routing, RIP routing, OSPF routing, ACLs, robust QoS, policy-based routing, low latency, and resiliency
- Scalable line-rate 40GbE for wireless traffic aggregation

Aruba 3810—The Aruba 3810M is available in a 16 port SFP+ and a two-module slot model. The module slots allow for an additional 8 SFP+ or 2 40-GbE ports. This switch is ideal for organizations with a small LAN who to aggregate 1 or 10-GbE connected access switches.

- Layer-3 switch with backplane stacking, static routing, RIP routing, OSPF routing, ACLs, robust QoS, policy-based routing, low latency, and resiliency
- Modular line-rate 10-GbE and 40-GbE ports for wireless aggregation

The next section of this guide helps you deploy the Mobile First Campus design in your organization.

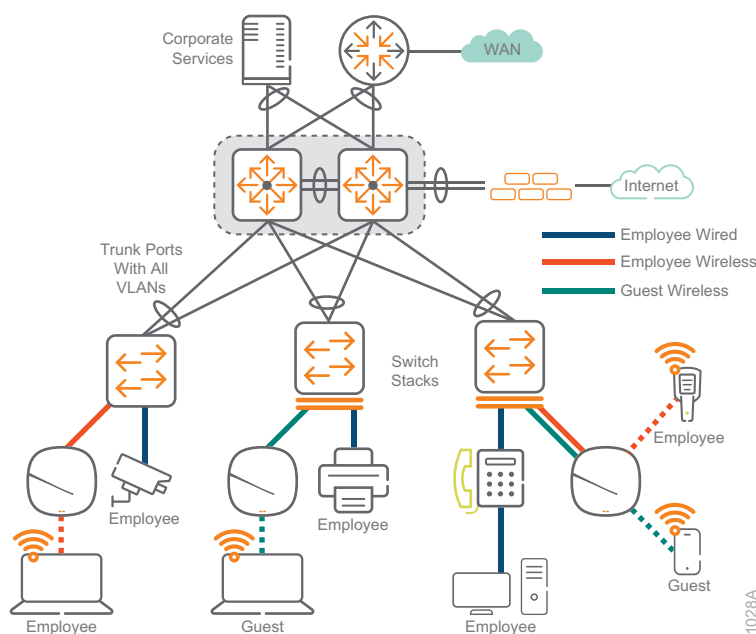
Deploying the Mobile First Campus

The Aruba Mobile First Campus design provides wired and wireless connectivity for local users. The wired LAN interconnects the wireless APs, WAN, data center, and Internet DMZ, making it a critical part of the network. Campus networks require a high-availability design to support mission-critical applications and real-time multimedia communications that drive organizational operations. The Mobile First Campus Design provides the following benefits:

- Specific functions of individual layers make the network easier to operate and maintain
- Modular building blocks quickly scale as the network grows
- Location-independent network access improves employee and guest productivity
- Hard-to-wire locations receive network connectivity without costly construction
- Plug-and-play wireless deployment with wired LAN switches preconfigured to recognize APs
- Centralized control of wireless environment is easy to manage and operate
- Reliable wireless connectivity, including complete radio frequency (RF) spectrum management is available with key Aruba management features

Simple, repeatable designs are easier to deploy, manage, and maintain. This design shows the most common and best supported options with general guidance for which option to choose. An overview of the Mobile First Campus design for 500 users is shown in the following diagram.

Figure 13 Mobile First Campus design overview



CAMPUS WIRED LAN

The wired LAN uses a hierarchical design model. Each layer performs specific functions helping to simplify the solution. In a typical network of up to 500 users, the wired LAN will have an aggregation layer and access layer. With the Aruba design, the trunks between the two layers use multiple active links forwarding traffic for a higher-performance network while reducing the complexity involved in traditional redundant designs.

Wired Access

The access layer in this design provides layer-2 connectivity to the network for wired and wireless devices. The layer-2 switches range from a single 2930F, 2930M, and 3810M to stacks of 2930Fs, 2930Ms and 3810Ms, along with a pair of stacked 5406. They are dual-connected to the dual-switch aggregation layer. Each uplink is connected to one of the two switches at the aggregation layer. If the access switches are stacked, the distributed ports are connected from different physical switches in the access layer.

The access switches are layer-2 and they contain four VLANs, one for management, employee wired, employee wireless, and guest wireless. For management purposes, each switch has an IP address in the management VLAN with a default gateway configured as the first-hop aggregation switch.

Procedures

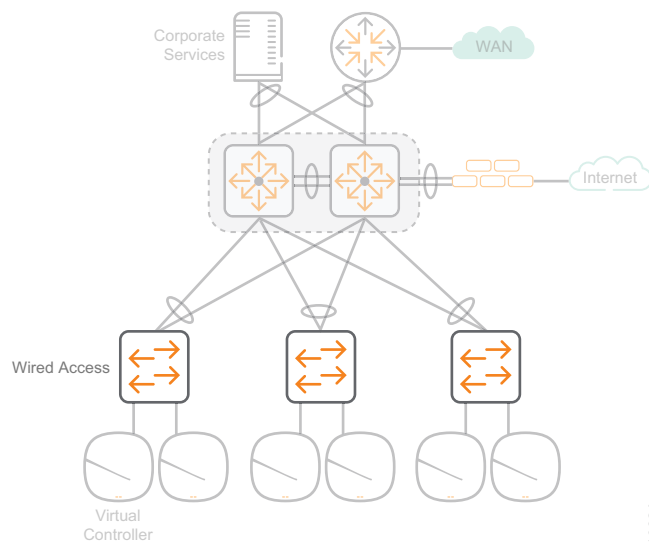
Configuring the Access Switch

- 1.1 Configure access switch stacking
- 1.2 Configure access-switch base features
- 1.3 Configure uplink ports from access to aggregation
- 1.4 Configure access-switch VLANs
- 1.5 Configure device profile for wireless access points
- 1.6 Configure the access-switch default gateway
- 1.7 Configure multicast IGMP snooping
- 1.8 Configure access-switch port-security features

Use this section for the access layer and repeat it for each wired access switch. This section can be used for stand-alone switches, switch stacks, or modular access switches.

The diagram below shows the wired access switch location in the Mobile First Campus design.

Figure 14 Mobile First Campus design—Wired access



1.1 Configure access switch stacking

Optional

This optional procedure is for switch platforms with backplane stacking modules using stack cables or front plane stacking using Virtual Switching Framework (VSF). If you are not using a switch stack in this area of your network, skip this procedure.

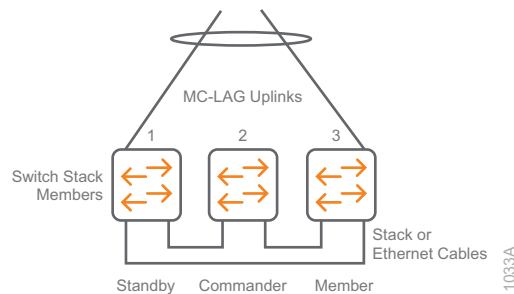
Stacking allows multiple access switches connected to each other through dedicated stacking ports or Ethernet connections to behave like a single switch. Stacking increases the port density by combining multiple physical devices into one virtual fabric, allowing management and configuration from one IP address. This reduces the total number of managed devices while better utilizing the port capacity in an access wiring closet. The members of a stack share the uplink ports, providing additional bandwidth and redundancy. There are three stacking-device roles:

- **Commander**—Conducts overall management of the stack, and manages the forwarding databases, synchronizing them with the standby.
- **Standby**—Provides redundancy for the stack and takes over stack management operations if the commander becomes unavailable or if an administrator forces a commander failover.
- **Members**—Are not part of the overall stack management; however, they must manage their local subsystems and ports to operate correctly as part of the stack. The commander and standby are also responsible for their own local subsystems and ports.

The device role is determined by member priority. When all switches in the stack are booted simultaneously, the switch with the highest priority becomes commander and the next highest priority becomes standby. The stacking priority can be set to any value between 1 and 255, and the default value is 128.

When connecting three or more switches into a logical switch stack, a ring topology is recommended. In a three-switch stack, connect switch one to switch two, connect switch two to switch three and connect switch three back to switch one to form a ring as shown in the diagram below. If a switch stack has three or more members, we recommend assigning the commander role to a switch that does not have uplinks to minimize forwarding delays when the commander becomes unavailable.

Figure 15 Three-switch ring topology and roles



If you are planning to use dedicated stacking modules with 2930M or 3810M switches, choose option 1. If you are planning to use Ethernet ports and VSF with 5400R or 2930F switches, choose option 2.

Option 1: Backplane stacking

The backplane stacking feature allows you to connect as many as ten switches into a single logical switch for data plane and management functions. One switch is designated as the commander, a second switch is configured as the standby, and other switches are designated as role member.

The following tables show the configuration details for backplane stacking.

Table 2 Backplane stacking for two-member switch stacks

	Switch 1	Switch 2
Stacking Member ID	1	2
Stacking Priority	230 (Standby)	250 (Commander)
Uplink	Yes	Yes

Table 3 Backplane stacking for three-member or more switch stacks

	Switch 1	Switch 2	Switch 3+
Stacking Member ID	1	2	3
Stacking Priority	230 (Standby)	250 (Commander)	128 (Member) default
Uplink	Yes	No	Yes

On a stack of three or more switches, assign the Commander role to a switch without uplinks. If your stack only has two switches, pick either switch for the Commander role because they will both have uplink ports.

Follow the steps below to connect the switches and statically assign their roles.

Step 1: Install the backplane stacking modules in all switches and connect the cables in a ring or mesh topology.

Step 2: Power-on each switch.

Note When the switches see each other through the stacking modules, stacking is enabled by default and member ID numbers will automatically be assigned.



Step 3: Display the member ID for each switch, using the **show stacking** command.

```
show stacking
```

```
...
```

ID	Mac Address	Model	Pri	Status
1	ecebb8-17f300	Aruba JL073A 3810M-24G-PoE+-1-slot...	128	Commander
*2	ecebb8-177480	Aruba JL073A 3810M-24G-PoE+-1-slot...	128	Member
3	ecebb8-175480	Aruba JL073A 3810M-24G-PoE+-1-slot...	128	Standby

Note The * indicates the physical switch you are using to view the stack.



Step 4: Following the guidelines in Table 1 and Table 2, determine the switch that will receive the Commander role and the switch that will receive the Standby role. If you have more than 2 switches in a stack, the additional switches will receive the Member role.

Step 5: On the stacking member that will receive the Commander role, configure the highest priority.

```
stacking member 2 priority 250
```

Step 6: On the stacking member that will receive the Standby role, configure the second highest priority.

```
stacking member 1 priority 230
```

Step 7: Save the configuration for all stack members.

```
write memory
```

Step 8: Reboot the switch stack for the changes to take effect.

```
boot system
```

```
This will reboot the system from the primary image.
```

```
Continue (y/n)? y
```

Step 9: After the switch stack reboots, verify stack status changes with the **show stacking** command.

```
show stacking
```

```
...
```

```

ID  Mac Address          Model                                     Pri Status
---
 1  ecebb8-17f300         Aruba JL073A 3810M-24G-PoE+-1-slot... 230 Standby
*2  ecebb8-177480         Aruba JL073A 3810M-24G-PoE+-1-slot... 250 Commander
 3  ecebb8-175480         Aruba JL073A 3810M-24G-PoE+-1-slot... 128 Member

```

Option 2: VSF stacking

VSF stacking allows switches to connect to each other through Ethernet ports in order to behave like a single logical switch. Like backplane stacking, the VSF fabric uses unique member IDs to identify and manage its members.

The VSF stack can have as many as four switches. The stack is formed using VSF links, which are logical interfaces comprised of same-speed physical interfaces. With the recommended ring topology, two logical VSF links are required per switch, one for each adjacent switch. For two-switch VSF stacks, only one logical VSF link is required.

The following tables show the configuration details for VSF stacking.

Table 4 VSF stacking for two-member switch stacks

	Switch 1	Switch 2
VSF Member	1	2
VSF Links	1	1
Priority	230 (Standby)	250 (Commander)
VSF Domain	200	200
Uplink	Yes	Yes

Table 5 VSF stacking for three-member or four-member switch stacks

	Switch 1	Switch 2	Switch 3+
VSF Member	1	2	3
VSF Links	1 and 2	1 and 2	1 and 2
Priority	230 (Standby)	250 (Commander)	128 (Member) default
VSF Domain	300	300	300
Uplink	Yes	No	Yes

On a stack of three or more switches, assign the Commander role to a switch without uplinks. If your stack only has two switches, pick either switch for the Commander role because they will both have uplink ports.

Follow the steps below to connect the switches and statically assign their roles in the stack.

Caution To prevent the half-configured links from causing problems, configure VSF prior to cabling the switches together.



Step 1: Following the guidelines in Table 3 and Table 4, determine the switch that will receive the Commander role and the switch that will receive the Standby role. If you have more than 2 switches in a stack, the additional switches will receive the Member role.

Step 2: On the switch that will receive the Standby role, configure the first member number ID with VSF link 1 and assign physical ports to it.

```
vsf member 1 link 1 A1-A2
```

Note To enable a VSF link, you must bind a minimum of one physical interface to it. The physical interfaces assigned to a VSF link automatically form an aggregate VSF link. A VSF link goes down only if all its VSF physical interfaces are down.



Step 3: For switches in a stack of three or more, configure the same member number ID with VSF link 2 and assign physical ports to it. Skip this step for two-member VSF switch stacks, because a second link is not needed.

```
vsf member 1 link 2 A3-A4
```

Step 4: Assign the Standby role to the switch, by configuring it with the second highest priority.

```
vsf member 1 priority 230
```

Step 5: Enable and save the configuration for the VSF domain.

```
vsf enable domain 300
```

This will save the current configuration and reboot the switch.

```
Continue (y/n)? y
```

Step 6: Connect to the switch that will receive the Commander role.

Step 7: Configure the member number ID with VSF link 1 and assign physical ports to it.

```
vsf member 2 link 1 A1-A2
```

Step 8: For switches in a stack of three or more, configure the same member number ID with VSF link 2 and assign physical ports to it. Skip this step for two-switch VSF stacks, because a second link is not needed.

```
vsf member 2 link 2 A3-A4
```

Step 9: Assign the Commander role to the switch, by configuring it with the highest priority.

```
vsf member 2 priority 250
```

Step 10: Enable and save the configuration for the VSF domain.

```
vsf enable domain 300
```

This will save the current configuration and reboot the switch.

```
Continue (y/n)? y
```

Step 11: If there are no additional switches, skip to Step 15.

Step 12: Connect to a switch that will receive the Member role.

Step 13: Configure the member with VSF links 1 and 2, and assign physical ports to the links.

```
vsf member 3 link 1 A1-A2
```

```
vsf member 3 link 2 A3-A4
```

Step 14: Enable and save the configuration for the VSF domain.

```
vsf enable domain 300
```

This will save the current configuration and reboot the switch.

```
Continue (y/n)? y
```

Step 15: For each additional switch in VSF stack, repeat Step 12 through Step 14, changing the variables according to the switch member ID and the physical ports assigned to the link.

Step 16: After all the switches in the stack are configured and rebooted, connect the VSF Ethernet ports.

Step 17: Use the following command to verify the VSF stack is operational.

```
show vsf topology
```

Example: Two-member VSF stack

VSF member's interconnection with links:

```
Stby      Cmdr
+---+    +---+
| 1 |1==1| 2 |
+---+    +---+
```

Example: Three-member VSF stack

VSF member's interconnection with links:

```
Stby      Cmdr
+---+    +---+    +---+
| 1 |1==2| 2 |1==2| 3 |
+---+    +---+    +---+
  2                      1
+=====+
```

1.2 Configure access-switch base features

In this procedure, you configure the base features for each access switch.

The switch has two levels of access: manager and operator. The manager has access to all areas of the configuration and has the ability to make configuration changes. The operator has access to the status, counters, and the event log, but the operator has read-only access to the command line interface and thus cannot make changes. You can only have one username and password for each level of access. The usernames are optional but we recommend changing them for additional security.

On each access switch, perform the following steps:

Step 1: Configure the switch host name.

```
hostname Access-Switch
```

Step 2: Configure the restricted operator username and password.

```
password operator user-name adminOper plaintext [passwordOper]
```

Step 3: Configure the unrestricted manager username and password

```
password manager user-name adminMgr plaintext [passwordMgr]
```

Step 4: Enable the SSH for inbound connections.

```
ip ssh
```

Step 5: Enable the secure copy protocol (SCP).

```
ip ssh filetransfer
```

Step 6: For increased security, turn off telnet server in order to only allow inbound SSH connections.

```
no telnet-server
```

Step 7: Configure a login banner.

```
banner motd #  
Property of example.com !! Unauthorized use prohibited !!  
#
```

Step 8: Configure the network time protocol (NTP) with time zone and daylight savings time. The time zone offset is entered as the difference in minutes from Coordinated Universal Time (UTC). The negative value means the amount of time behind UTC. The NTP iburst feature provides faster time synchronization.

```
time daylight-time-rule continental-us-and-canada  
time timezone -480  
timesync ntp  
ntp unicast  
ntp server 10.2.120.70 iburst  
ntp enable
```

Step 9: If the date on your device is not current, use the time command to set the date to today's date. The current date is required so that, in the next step, you can create a valid certificate.

```
time MM/DD/YYYY
```

Step 10: Configure HTTP Secure (HTTPS) for web access to the switch.

```
crypto pki identity-profile https_Profile subject  
Enter Common Name(CN) : ExampleSwitch  
Enter Org Unit(OU) : ExampleOrgUnit  
Enter Org Name(O) : ExampleOrg  
Enter Locality(L) : Roseville  
Enter State(ST) : California  
Enter Country(C) : US  
crypto pki enroll-self-signed certificate-name https_Certificate  
web-management ssl
```

Step 11: For additional security, turn off plaintext HTTP management.

```
no web-management plaintext
```

Step 12: Enable the simple network management protocol version 3 (SNMPv3).

```
snmpv3 enable
SNMPv3 Initialization process.
Creating user 'initial'
Authentication Protocol: MD5
Enter authentication password: [password]
Privacy protocol is DES
Enter privacy password: [password]

User 'initial' has been created
Would you like to create a user that uses SHA? [y/n] y

User creation is done.  SNMPv3 is now functional.
Would you like to restrict SNMPv1 and SNMPv2c messages to have read only
access (you can set this later by the command 'snmpv3 restricted-access')?
[y/n] n
```

Step 13: Create full read-write, limited read-write and read-only users for SNMPv3.

```
snmpv3 user NetAdminRW auth sha [passwordRW] priv aes [passwordRW]
snmpv3 user NetAdminLimited auth sha [passwordLimited] priv aes [passwordLim-
ited]
snmpv3 user NetAdminR auth sha [passwordRO] priv aes [passwordRO]
```

Step 14: For additional security, remove the SNMP server community public from the configuration.

```
no snmp-server community public
```

1.3 Configure uplink ports from access to aggregation

The uplink ports use the link aggregation control protocol (LACP) to combine two or more physical ports into a single trunk interface. By default, the uplink trunks will use source and destination IP addresses to load-balance traffic between the physical interfaces. If a VLAN is not specified in the link-keepalive command, the unidirectional link detection (UDLD) packets are sent untagged.

Step 1: Configure the dual-port trunks with LACP.

```
trunk 2/A1,3/A1 trk11 lacp
```

Step 2: Configure UDLD on the uplink ports, set the interval to 70 (70 at 100-ms increments = 7 seconds) and the retries to 6.

```
interface 2/A1,3/A1 link-keepalive
link-keepalive interval 70
link-keepalive retries 6
```

Step 3: Enable Spanning Tree Protocol (STP) globally on the switch.

```
spanning-tree enable
```

Step 4: Increase the logging level to informational, for visibility to additional link and trunk status events.

```
logging severity info
```

1.4 Configure access-switch VLANs

The layer-2 access switches need an IP address on the management VLAN, for operational purposes. The non-trunk ports are configured as untagged in the wired VLAN. The trunk ports are configured as tagged for the user VLANs and untagged for VLAN 777.

VLAN hopping is a computer security exploit that uses double tagging to attack network resources on a VLAN. The basic concept behind all VLAN hopping attacks is for an attacking host on a VLAN to gain access to traffic on other VLANs that would normally not be accessible.

Double tagging can be mitigated by creating an unused VLAN that will only be configured as the native VLAN on uplink trunk interfaces. The unused VLAN 777 does not have an IP address and it is not connected to anything else on the switch.

The following table provides the VLAN assignments for the Mobile First Campus design.

Table 6 VLAN assignments, IP subnets, and port tagging

VLAN name	VLAN ID	IP address	Tagged/Untagged ports
Management	10	10.2.0.10/22	Tagged Trk11
Wired	20	N/A	Untagged 1/1-1/48 (all non-trunk ports) Tagged Trk11
Wireless	30	N/A	Tagged Trk11
Guest	40	N/A	Tagged Trk11
Anti-VLAN hopping	777	N/A	Untagged Trk11

On each access switch, perform the following steps:

Step 1: For each VLAN in Table 5, configure the VLAN.

Example: Management VLAN

```
vlan 10
  name Management
  tagged Trk11
  ip address 10.2.0.10 255.255.252.0
  exit
```

Example: Anti-VLAN hopping VLAN with no IP address

```
vlan 777
  name Anti-VLAN hopping
  untagged Trk11
  exit
```

Step 2: Enable Rapid Per-VLAN Spanning Tree protocol (Rapid-PVST).

```
spanning-tree mode rapid-pvst
```

Step 3: Use the management VLAN IP address to configure the source address for SNMP responses from the switch.

```
snmp-server response-source 10.2.0.10
```

1.5 Configure device profile for wireless access points

In this procedure, the access VLANs you previously configured for management, employee, and guest traffic are added to the device profile that the switches will apply to traffic received from a connected AP.

The device profile in this design is used to apply the untagged and tagged VLAN commands to the port where the AP is connected. The untagged VLAN is used by the AP to communicate with other APs and the virtual controller. The tagged VLANs allow employee and guest wireless traffic to remain segmented by the switch.

On each access switch, perform the following steps:

Step 1: Configure the device profile name.

```
device-profile name "Aruba-AP-Profile"  
  untagged-vlan 10  
  tagged-vlan 30,40
```

Step 2: Configure the device profile type.

```
device-profile type "aruba-ap"  
  associate "Aruba-AP-Profile"  
  enable
```

1.6 Configure the access-switch default gateway

The IP default gateway is necessary to forward traffic sourced from the switch to the management VLAN and the rest of the network, using the IP address of the aggregation switch as its next hop router.

On each access switch, perform the following step:

Step 1: Configure the IP default gateway for the management VLAN.

```
ip default-gateway 10.2.0.1
```

1.7 Configure multicast IGMP snooping

This procedure enables multicast IGMP snooping for the layer-2 access switches.

On each access switch, perform the following step:

Step 1: Configure multicast IGMP snooping on the employee wired and wireless VLANs.

```
vlan 20 ip igmp  
vlan 30 ip igmp
```

1.8 Configure access-switch port-security features

This procedure configures port security for the access switches. DHCP snooping for IPv4 and IPv6 stops DHCP starvation attacks and it also prevents rogue DHCP servers from servicing requests on your network. ARP protect stops man-in-the-middle attacks caused by ARP cache poisoning. Dynamic IP lockdown stops devices from forging

their source IP address by inspecting the IP-MAC binding information in the DHCP snooping table. IPv6 RA guard stops rogue IPv6 clients from advertising themselves as routers. BPDU protection prevents loops in the network by putting a non-trunk port into a disabled state for a specified amount of time when it receives a BPDU from another switch.

Caution Although these features are recommended for a secure access layer, they should be applied after the network is fully operational, in order to avoid problems during the initial stages of building the network.



Apply the features one at a time and check the logs if connectivity problems begin.

Step 1: Enable DHCP snooping and configure it on all VLANs and trust the trunk interface.

```
dhcp-snooping
dhcp-snooping vlan 10 20 30 40 777
dhcp-snooping trust trk11
```

Step 2: Enable DHCPv6 snooping and configure it on all VLANs and trust the trunk interface.

```
dhcpv6-snooping
dhcpv6-snooping vlan 10 20 30 40 777
dhcpv6-snooping trust trk11
```

Step 3: Enable ARP protection and configure it on all VLANs, except the management VLAN 10 and trust the trunk interface.

```
arp-protect
arp-protect vlan 20 30 40 777
arp-protect trust trk11
```

Step 4: Enable IP source guard globally.

```
ip source-lockdown
```

Step 5: Configure IPv6 router advertisement (RA) guard on the range of non-trunk ports.

```
ipv6 ra-guard ports ethernet 1/1-1/48
```

Step 6: (Optional) Configure spanning tree BPDU protection on the range of non-trunk ports and configure the port to be disabled for 60 seconds.

Caution This command will shut down a port for 60 seconds if a device that sends BPDUs is connected. Certain IP phones with built-in switches send BPDUs, so you will have to trust ports with these types of devices.



```
spanning-tree 1/1-1/48 bpdu-protection
spanning-tree bpdu-protection-timeout 60
```

Wired Aggregation

The aggregation layer provides connectivity for all the access layer switches and connects to external networks in the campus LAN.

Procedures

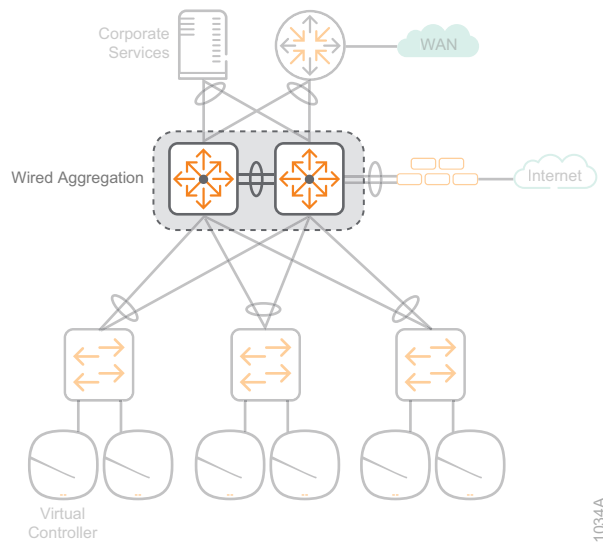
Configuring the Aggregation Switch

- 2.1 Configure aggregation-switch stacking
- 2.2 Configure the aggregation-switch base features
- 2.3 Configure uplink ports from aggregation to access
- 2.4 Configure aggregation-switch VLANs
- 2.5 Configure OSPF routing
- 2.6 Configure IP multicast routing

Use this section for the aggregation layer and repeat it for each wired aggregation switch. This section can be used for standalone switches, switch stacks or modular aggregation switches.

The diagram below shows the wired aggregation switch location in the Mobile First Campus design.

Figure 16 Mobile First Campus design—Wired aggregation



2.1 Configure aggregation-switch stacking

Optional

This optional procedure is for switch platforms with backplane stacking modules using stack cables or front plane stacking using VSF. Skip this procedure if you are not using a switch stack in this area of your network.

Stacking allows multiple access switches connected to each other through dedicated stacking ports or Ethernet connections to behave like a single switch. Stacking increases the port density by combining multiple physical devices into one virtual fabric, allowing management and configuration from one IP address. The members of a stack share the uplink ports providing additional bandwidth and redundancy.

There are three stacking device roles:

- **Commander**—Conducts overall management of the stack, and manages the forwarding databases, synchronizing them with the standby.
- **Standby**—Provides redundancy for the stack and takes over stack management operations if the commander becomes unavailable, or if an administrator forces a commander failover.
- **Members**—Are not part of the overall stack management; however, they must manage their local subsystems and ports to operate correctly as part of the stack. The commander and standby are also responsible for their own local subsystems and ports.

The device role is determined by member priority. When all switches in the stack are booted simultaneously, the switch with the highest priority becomes commander and the next highest priority becomes standby. The stacking priority can be set to any value between 1 and 255, and the default value is 128.

In this design, we recommend a maximum of two switches in the stack for the aggregation layer. Smaller networks can use two 3810M switches and larger networks can use two 5400R switches.

If you are planning to use dedicated stacking modules with 3810M switches, choose option 1. If you are planning to use Ethernet ports and VSF with 5400R switches, choose option 2.

Option 1: Backplane stacking

The backplane stacking feature allows you to connect as many as ten switches into a single logical switch for data plane and management functions. In the aggregation layer, we recommend only using two switches. One switch is designated as the commander and the second switch is configured in the standby role.

The table below shows the configuration details for backplane stacking.

Table 7 Backplane stacking for two-member switch stacks

	Switch 1	Switch 2
Stacking Member ID	1	2
Stacking Priority	230 (Standby)	250 (Commander)
Uplink	Yes	Yes

Follow the steps below to connect the switches and statically assign their roles.

Step 1: Install the backplane stacking modules in all switches and connect the cables.

Step 2: Power-on each switch.

Step 3: Display the member ID for each switch using the **show stacking** command.

```
show stacking
```

```
...
```

```
ID  Mac Address          Model                                     Pri Status
---  -
  1  9457a5-8c3080         Aruba JL075A 3810M-16SFP+-2-slot S... 128 Commander
 *2  9457a5-8c9000         Aruba JL075A 3810M-16SFP+-2-slot S... 128 Standby
```

Step 4: Assign the Commander role to a switch, by configuring the switch to have the highest priority.

```
stacking member 2 priority 250
```

Step 5: Assign the Standby role to the other switch, by configuring the switch to have the second highest priority.

```
stacking member 1 priority 230
```

Step 6: Save the configuration for all stack members.

```
write memory
```

Step 7: Reboot the switch stack for the changes to take effect.

```
boot system
```

```
This will reboot the system from the primary image.
```

```
Continue (y/n)? y
```

Step 8: After the switch stack reboots, verify stack status changes with the **show stacking** command.

```
show stacking
```

```
...
```

ID	Mac Address	Model	Pri	Status
1	9457a5-8c3080	Aruba JL075A 3810M-16SFP+-2-slot S...	230	Standby
*2	9457a5-8c9000	Aruba JL075A 3810M-16SFP+-2-slot S...	250	Commander

Option 2: VSF Stacking

VSF stacking allows switches to connect to each other through Ethernet ports in order to behave like a single logical switch. Like backplane stacking, the VSF fabric uses unique member IDs to identify and manage its members.

The VSF stack is formed using VSF links, which are logical interfaces comprised of same-speed physical interfaces. For two-member VSF switch stacks, only one logical VSF link is required.

In the aggregation layer, we recommend only using two switches. One switch is designated as the Commander and the second switch is configured in the Standby role.

The table below shows the configuration details for VSF stacking.

Table 8 VSF stacking for two-member switch stacks

	Switch 1	Switch 2
VSF Member	1	2
VSF Links	1	1
Priority	230 (Standby)	250 (Commander)
VSF Domain	200	200
Uplink	Yes	Yes

Follow the steps below to connect the switches and statically assign their roles in the stack.

Step 1: Configure the first member number ID with VSF link 1 and assign physical ports to it.

```
vsf member 1 link 1 A1-A2
```

Step 2: Assign the Standby role to the switch, by configuring it with the second highest priority.

```
vsf member 1 priority 230
```

Step 3: Enable and configure VSF domain.

```
vsf enable domain 200
```

This will save the current configuration and reboot the switch.

```
Continue (y/n)? y
```

Step 4: Connect to the second switch.

Step 5: Configure the second member number ID with VSF link 1 and assign physical ports to it.

```
vsf member 2 link 1 A1-A2
```

Step 6: Assign the Commander role to the switch, by configuring it with the highest priority.

```
vsf member 2 priority 250
```

Step 7: Enable and configure VSF domain.

```
vsf enable domain 200
```

This will save the current configuration and reboot the switch.

```
Continue (y/n)? y
```

Step 8: After both switches in the stack are configured and rebooted, connect the VSF Ethernet ports.

Step 9: Use the following command to verify the VSF stack is operational.

```
show vsf topology
```

Example: Two-member VSF stack

VSF member's interconnection with links:

```
Stby      Cmdr
+---+      +---+
| 1 |1==1| 2 |
+---+      +---+
```

2.2 Configure the aggregation-switch base features

In this procedure, you configure the base features for each aggregation switch.

The switch has two levels of access: manager and operator. The manager has access to all areas of the configuration and has the ability to make changes. The operator has access to the status, counters, and the event log, but the operator has read-only access to the command line interface and thus cannot make changes. You can only have one username and password for each level of access. The usernames are optional, but we recommend changing them for additional security.

On each aggregation switch, perform the following steps:

Step 1: Configure the switch host name.

```
hostname Aggregation-Switch
```

Step 2: Configure the restricted operator username and password.

```
password operator user-name adminOper plaintext [passwordOper]
```

Step 3: Configure the unrestricted manager username and password

```
password manager user-name adminMgr plaintext [passwordMgr]
```

Step 4: Enable SSH for inbound connections.

```
ip ssh
```

Step 5: Enable SCP.

```
ip ssh filetransfer
```

Step 6: For increased security, turn off telnet server in order to only allow inbound SSH connections.

```
no telnet-server
```

Step 7: Configure a login banner.

```
banner motd #  
Property of example.com !! Unauthorized use prohibited !!  
#
```

Step 8: Configure the network time protocol (NTP) with time zone and daylight savings time. The iburst feature provides faster time synchronization. The time zone offset is entered as the difference in minutes from Coordinated Universal Time (UTC). The negative value means the amount of time behind UTC.

```
time daylight-time-rule continental-us-and-canada  
time timezone -480  
timesync ntp  
ntp unicast  
ntp server 10.2.120.70 iburst  
ntp enable
```

Step 9: If the date on your device is not current, use the time command to set the date to today's date. The current date is required so that, in the next step, you can create a valid certificate.

```
time MM/DD/YYYY
```

Step 10: Configure HTTPS for web access to the switch.

```
crypto pki identity-profile https_Profile subject  
Enter Common Name(CN) : ExampleSwitch  
Enter Org Unit(OU) : ExampleOrgUnit  
Enter Org Name(O) : ExampleOrg  
Enter Locality(L) : Roseville  
Enter State(ST) : California  
Enter Country(C) : US  
crypto pki enroll-self-signed certificate-name https_Certificate  
web-management ssl
```

Step 11: For additional security, turn off plaintext HTTP management.

```
no web-management plaintext
```

Step 12: Enable the simple network management protocol version 3 (SNMPv3).

```
snmpv3 enable
SNMPv3 Initialization process.
Creating user 'initial'
Authentication Protocol: MD5
Enter authentication password: [password]
Privacy protocol is DES
Enter privacy password: [password]

User 'initial' has been created
Would you like to create a user that uses SHA? [y/n] y

User creation is done.  SNMPv3 is now functional.
Would you like to restrict SNMPv1 and SNMPv2c messages to have read only
access (you can set this later by the command 'snmpv3 restricted-access')?
[y/n] n
```

Step 13: Create full read-write, limited read-write and read-only users for SNMPv3.

```
snmpv3 user NetAdminRW auth sha [passwordRW] priv aes [passwordRW]
snmpv3 user NetAdminLimited auth sha [passwordLimited] priv aes [passwordLim-
ited]
snmpv3 user NetAdminR auth sha [passwordRO] priv aes [passwordRO]
```

Step 14: For additional security, remove the SNMP server community public from the configuration.

```
no snmp-server community public
```

2.3 Configure uplink ports from aggregation to access

The uplink ports use LACP to combine two or more physical ports into a single trunk interface. By default, the uplink trunks will use source and destination IP addresses to load balance traffic between the physical interfaces. If a VLAN is not specified in the link-keepalive command, the UDLD packets are sent untagged

On each aggregation switch, perform the following steps:

Step 1: Configure the dual-port trunks with LACP.

```
trunk 1/A1,2/A1 trk11 lacp
```

Repeat this step for each trunk.

Step 2: Configure UDLD on the uplink ports.

```
int 1/A1,2/A1 link-keepalive
```

Repeat this step for each set of uplink ports.

Step 3: Configure the keepalive interval to 70 (70 at 100-ms increments = 7 seconds) and the retries to 6.

```
link-keepalive interval 70
```

```
link-keepalive retries 6
```

Step 4: Enable STP globally on the switch. Configure the spanning tree priority to 0, which is the highest priority and makes the aggregation switch the spanning tree root bridge.

Note A root bridge should always be statically defined to prevent a rogue or mis-configured switch from altering the STP topology.



```
spanning-tree enable
```

```
spanning-tree priority 0
```

Step 5: Increase the logging level to informational for visibility to additional link and trunk status events.

```
logging severity info
```

2.4 Configure aggregation-switch VLANs

The layer-3 aggregation switch is the default gateway for the user VLANs and they need an IP address. The non-trunk ports are configured as untagged in the wired VLAN. The uplink trunk ports are configured as tagged for the user VLANs and untagged for VLAN 777.

VLAN hopping is a computer security exploit that uses double tagging to attack network resources on a VLAN. The basic concept behind all VLAN hopping attacks is for an attacking host on a VLAN to gain access to traffic on other VLANs that would normally not be accessible.

Double tagging can be mitigated by creating an unused VLAN that will only be configured as the native VLAN on uplink trunk interfaces. The unused VLAN 777 does not have an IP address and it is not connected to anything else on the switch.

When you are using a centralized DHCP server, the **ip helper-address** command allows remote DHCP servers to provide end-station IP addresses for the VLAN. The helper command points to the IP address of the central DHCP server. If you have more than one DHCP server servicing the same VLAN, you can list multiple helper commands on an interface. The DHCP client will accept the first offer it receives.

The following table provides the VLAN assignments for the Mobile First Campus design.

Table 9 VLAN assignments, IP addresses, and tagging

VLAN name	VLAN ID	IP address	IP helper address	Tagged/Untagged ports
Management	10	10.2.0.1/22	10.2.120.40	Tagged Trk11-Trk15
Wired	20	10.2.4.1/22	10.2.120.40	Untagged 1/1-1/48 (all non-trunk ports) Tagged Trk11-Trk15
Wireless	30	10.2.8.1/22	10.2.120.40	Tagged Trk11-Trk15
Guest	40	10.2.12.1/22	10.2.120.40	Tagged Trk11-Trk15
Anti-VLAN hopping	777	N/A	N/A	Untagged Trk11-Trk15

On each aggregation switch, perform the following steps:

Step 1: Configure the aggregation VLANs.

Example: Management VLAN

```
vlan 10
  name "Management"
  tagged Trk11-Trk15
  ip address 10.2.0.1 255.255.252.0
  ip helper-address 10.2.120.40
exit
```

Example: Anti-VLAN hopping VLAN with no IP address and no helper address

```
vlan 777
  name Anti-VLAN hopping
  untagged Trk11-Trk15
exit
```

Repeat this step for each VLAN in the previous table.

Step 2: Enable Rapid-PVST. Configure the spanning tree priority on the access VLANs to 0, which is the highest priority and makes the aggregation switch the spanning tree root bridge.

```
spanning-tree mode rapid-pvst
spanning-tree vlan 10,20,30,40,777 priority 0
```

Step 3: Use the management VLAN IP address to configure the source address for SNMP responses from the switch.

```
snmp-server response-source 10.2.0.1
```

2.5 Configure OSPF routing

This procedure configures OSPF as the layer-3 routing protocol and it uses area backbone for the entire network. Redistribute the connected and static routes, restrict the routes to the 10.2.0.0/16 network and enable nonstop forwarding. Configure the user VLANs as passive because there will be no devices that need routing protocol updates attached to the user VLANs.

On each aggregation switch, perform the following steps:

Step 1: Configure the loopback interface.

```
interface loopback 1
ip address 10.2.255.10
```

Step 2: Enable IP routing. Configure the router ID as the IP address of the loopback interface from the previous step.

```
ip routing
ip router-id 10.2.255.10
```

Step 3: Configure OSPF.

```
router ospf
area backbone
redistribute connected
redistribute static
restrict 10.2.0.0 255.255.0.0
nonstop
enable
exit
```

Step 4: Configure the loopback interface for OSPF.

```
interface loopback 1
ip ospf 10.2.255.10 area backbone
```

Step 5: Configure the user VLANs for OSPF.

```
vlan 10
  ip ospf 10.2.0.1 passive
  ip ospf 10.2.0.1 area backbone
exit
```

Step 6: For each user VLAN with an IP address, repeat Step 5.

2.6 Configure IP multicast routing

This procedure enables multicast routing for the layer-3 aggregation switches. The design is based on sparse mode multicast operation. Bootstrap routers (BSRs) and rendezvous points (RPs) are used to provide a simple yet scalable way to provide a highly resilient RP environment.

The BSR priority range is from 0-255 and the default is 0. The candidate with the *highest* value becomes the BSR for the domain.

The RP priority range is from 0-255 and the default is 192. The candidate with the *lowest* value becomes the RP for the defined group of multicast prefixes.

On each aggregation switch, perform the following steps:

Step 1: Enable IP multicast routing in global configuration mode.

```
ip multicast-routing
```

Step 2: Enable PIM and configure the switch as a BSR candidate with a source VLAN 20 and priority of 50.

```
router pim
  enable
  bsr-candidate
  bsr-candidate source-ip-vlan 20
  bsr-candidate priority 50
```

Step 3: Configure the switch as a candidate RP with a source VLAN 20 and a group prefix of 224.0.0.0 to 240.0.0.0. Set the RP candidate hold time to 150 seconds and the priority to 50.

```
rp-candidate source-ip-vlan 20
rp-candidate group-prefix 224.0.0.0 240.0.0.0
rp-candidate hold-time 150 priority 50
```

Step 4: Configure PIM sparse mode on the employee VLANs and allow any IP address to source multicast streams.

```
vlan 20
  ip pim-sparse
  ip-addr any
  exit
exit
```

Step 5: For each of your employee VLANs where you want to send multicast traffic, repeat Step 4.

CAMPUS WIRELESS LAN

The WLAN provides network access for employees, wireless Internet access for guests, and connectivity for IoT devices. Regardless of their location on the network, wireless devices have the same experience when connecting to their services.

The wireless configuration consists of Instant APs (IAPs) with management, employee, and guest SSIDs. The APs use the management VLAN to communicate between each other and the virtual controller. The employee traffic is sent to the employee VLAN and the guest traffic to the guest VLAN. An access-rule policy created in the virtual controller allows the guests to access the DHCP server, the DNS service, and HTTP/HTTPS, in order to access web sites on the Internet.

Procedures

Configuring the Instant Access Point Virtual Controller

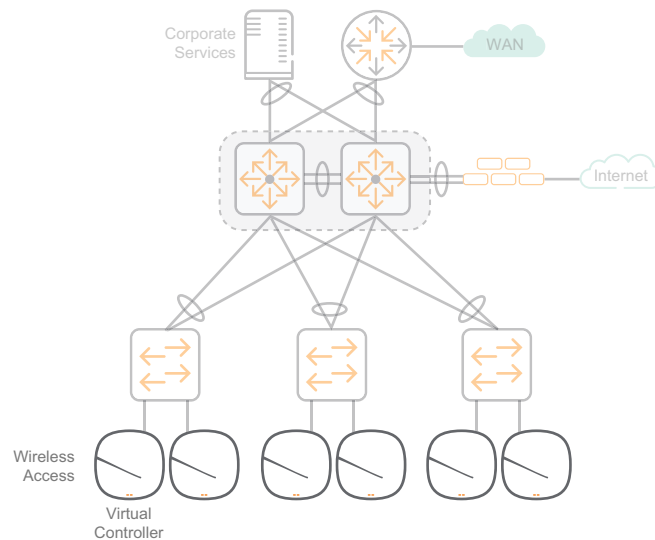
- 3.1 Configure the virtual controller system setup
- 3.2 Configure the radio frequency
- 3.3 Configure the power settings
- 3.4 Configure the employee SSID name and VLAN
- 3.5 Configure the employee SSID security
- 3.6 Configure the employee SSID access rules
- 3.7 Configure the guest SSID name and VLAN
- 3.8 Configure the guest SSID security
- 3.9 Configure the guest SSID access rules
- 3.10 Configure the access point name
- 3.11 Remove the SetMeUp SSID

Use this section for the wireless access layer. This section can be used for IAP virtual controllers only.

The cluster of IAPs are controlled by a single IAP that serves a dual role as virtual controller and IAP, which eliminates the need for dedicated hardware controllers. The virtual controller is elected from the cluster of IAPs in each layer-2 domain. Only the IAP acting as the virtual controller requires a configuration because the rest of them will inherit the necessary information from the virtual controller. The Instant operating system continually monitors the network to determine which IAP should act as the virtual controller, and the functionality will move from IAP to IAP as needed, without impacting performance.

The diagram below shows the wireless IAP cluster in the Mobile First Campus design, with an elected virtual controller in the layer-2 domain.

Figure 17 Mobile First Campus design—Wireless access



3.1 Configure the virtual controller system setup

This procedure configures the system setup for the virtual controller. After the wired network is operational, connect the first IAP to an access switch in a location you can easily reach. The first IAP in a layer-2 domain automatically becomes the virtual controller. The initial power-on sequence takes several minutes to complete.

Step 1: Connect the first IAP to a PoE port on an access switch.

Step 2: When the Radio Status light is blinking green, from your wireless PC connect to the open SSID that has the name "SetMeUp-XX:XX:XX".

Note Connecting to the SSID will automatically open your default web browser, but you should get a security warning saying the site is not secure.



Step 3: In the web browser that opens, click the option to proceed to the webpage. The following screenshot shows an example of the message you will see in your browser. Based on your browser type, you might see a slightly different message.

This site is not secure

This might mean that someone's trying to fool you or steal any info you send to the server. You should close this site immediately.



[Close this tab](#)



[More information](#)

Your PC doesn't trust this website's security certificate.

Error Code: DLG_FLAGS_INVALID_CA



[Go on to the webpage \(not recommended\)](#)

Note If your browser does not allow you to proceed to the web page due to security settings, you may have to use a different browser.

The other option is to browse directly to the DHCP IP address on the uplink port of the IAP.

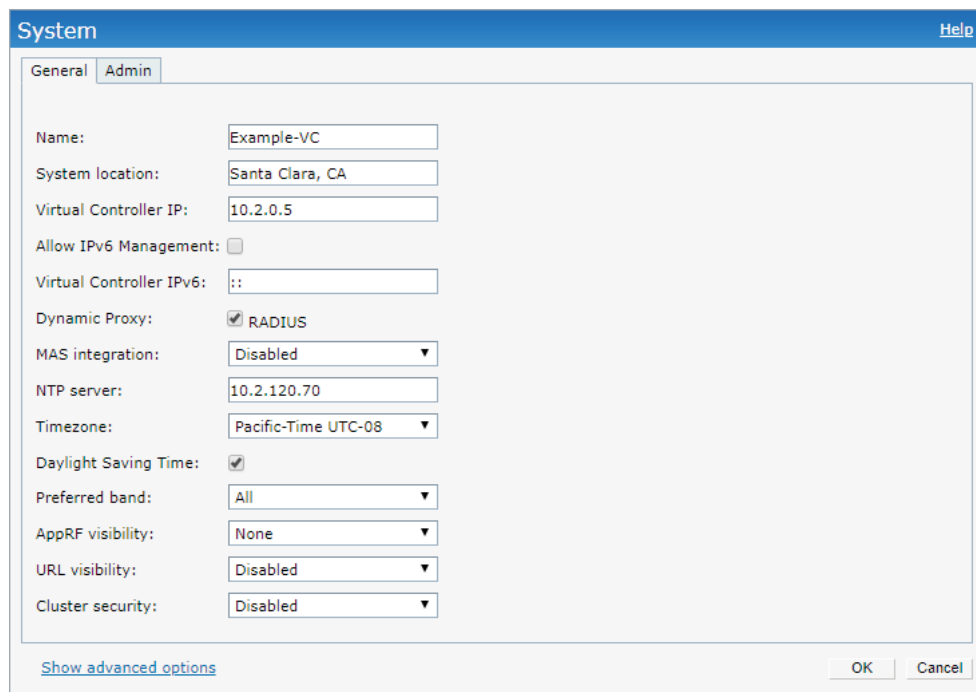


Step 4: On the **Virtual Controller** welcome page, enter the following information, and then click **Log In**.

- Username—**admin**
- Password—**admin**

Step 5: Navigate to **System > General**, enter the following information, and then click **OK**.

- Name—**Example-VC**
- System location—**Santa Clara, CA**
- Virtual Controller IP—**10.2.0.5** (use an IP address outside the DHCP scope)
- Dynamic Proxy—**Checkbox** (use the VC IP above for communication to RADIUS server)
- NTP Server—**10.2.120.70**
- Timezone—**Pacific-Time UTC-08**
- Daylight Saving Time—**Checkbox**



The screenshot shows the 'System' configuration window with the 'General' tab selected. The 'Admin' tab is also visible. The configuration fields are as follows:

Field	Value
Name	Example-VC
System location	Santa Clara, CA
Virtual Controller IP	10.2.0.5
Allow IPv6 Management	<input type="checkbox"/>
Virtual Controller IPv6	::
Dynamic Proxy	<input checked="" type="checkbox"/> RADIUS
MAS integration	Disabled
NTP server	10.2.120.70
Timezone	Pacific-Time UTC-08
Daylight Saving Time	<input checked="" type="checkbox"/>
Preferred band	All
AppRF visibility	None
URL visibility	Disabled
Cluster security	Disabled

At the bottom, there is a link for 'Show advanced options' and buttons for 'OK' and 'Cancel'.

Step 6: Navigate to **System > Admin**, and then enter the following information.

- Authentication—**Authentication server w/fallback to Internal**
- Auth server 1—**New**

Step 7: In the **New Server** popup window, enter the following information, and then click **OK**.

- Click the **RADIUS** radio button
- Name—**Example-Radius**

- IP address—**10.2.120.40** (IP address of RADIUS server)
- Auth port—**1812** (default)
- Accounting port—**1813** (default)
- Shared key—**[password]**
- Retype key—**[password]**
- RFC 3576—**Enabled**
- Air Group CoA port—**5999** (default)
- NAS IP address—**10.2.0.5** (IP address of this virtual controller from previous step)

New Server

☒ **RADIUS** ☐ **TACACS**

Name:

IP address:

RadSec:

Auth port:

Accounting port:

Shared key:

Retype key:

Timeout: sec.

Retry count:

RFC 3576:

Air Group CoA port:

RFC 5997: ☐ Authentication ☐ Accounting

NAS IP address: (optional)

NAS identifier: (optional)

Dead time: min.

DRP IP:

DRP Mask:

DRP VLAN:

DRP Gateway:

Service type framed user: ☐ 802.1X ☐ Captive Portal ☐ MAC

Step 8: After returning to the **System > Admin** page, change the default admin password by entering the following information, and then click **OK**.

- Password—[password]
- Retype—[password]

The screenshot shows the 'System > Admin' configuration window. The 'Admin' tab is active. The 'Local' section contains fields for 'Authentication' (set to 'Authentication server w/fallback to Internal'), 'Auth server 1' (set to 'Example-Radius' with an 'Edit' button), 'Auth server 2' (set to '-- Select Server --'), 'Username' (set to 'admin'), 'Password' (masked with dots), and 'Retype' (masked with dots). The 'View Only' section has an 'Authentication' dropdown set to 'Authentication serve'. The 'Guest Registration Only' section also has an 'Authentication' dropdown set to 'Authentication serve'. On the right, the 'AirWave' section includes fields for 'Organization', 'Airwave server', 'AirWave backup server', 'Shared key', and 'Retype'. At the bottom left, there is a link 'Show advanced options'. At the bottom right, there are 'OK' and 'Cancel' buttons.

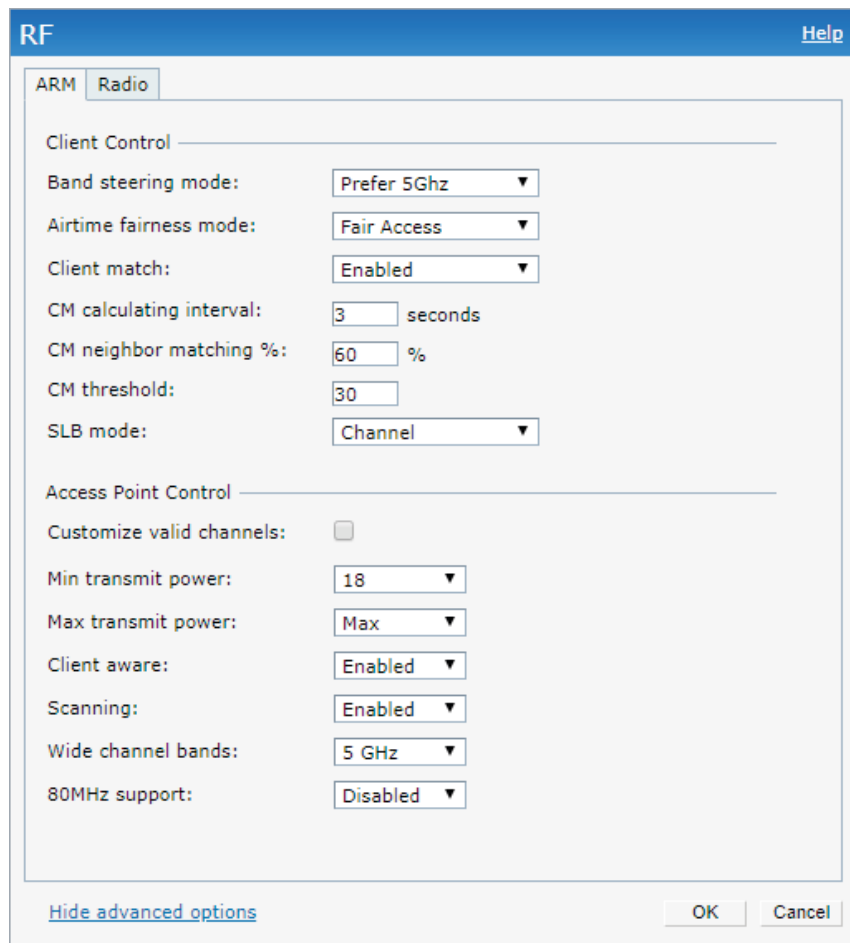
3.2 Configure the radio frequency

In this procedure, you configure the radio frequency for the 5-GHz and 2.4-GHz radios. Select fair access for the airtime fairness mode and enable ClientMatch. Change the ClientMatch threshold to 30 devices and disable support for the 80-MHz-wide and 160-MHz-wide channels.

Step 1: Navigate to **RF > ARM** and click **Show advanced options**.

Step 2: On the **Advanced options** page, enter the following information.

- Airtime fairness mode—**Fair Access**
- Client match—**Enabled**
- CM threshold—**30**
- 80MHz support—**Disabled**



The screenshot shows the 'RF' configuration window with the 'Radio' tab selected. The window has a blue header with 'RF' and a 'Help' link. Below the header are two tabs: 'ARM' and 'Radio'. The 'Radio' tab contains two sections: 'Client Control' and 'Access Point Control'. The 'Client Control' section includes: 'Band steering mode' (dropdown: Prefer 5Ghz), 'Airtime fairness mode' (dropdown: Fair Access), 'Client match' (dropdown: Enabled), 'CM calculating interval' (input: 3 seconds), 'CM neighbor matching %' (input: 60 %), 'CM threshold' (input: 30), and 'SLB mode' (dropdown: Channel). The 'Access Point Control' section includes: 'Customize valid channels' (checkbox: unchecked), 'Min transmit power' (dropdown: 18), 'Max transmit power' (dropdown: Max), 'Client aware' (dropdown: Enabled), 'Scanning' (dropdown: Enabled), 'Wide channel bands' (dropdown: 5 GHz), and '80MHz support' (dropdown: Disabled). At the bottom of the window are a 'Hide advanced options' link, an 'OK' button, and a 'Cancel' button.

Section	Parameter	Value
Client Control	Band steering mode	Prefer 5Ghz
	Airtime fairness mode	Fair Access
	Client match	Enabled
	CM calculating interval	3 seconds
	CM neighbor matching %	60 %
	CM threshold	30
	SLB mode	Channel
Access Point Control	Customize valid channels	<input type="checkbox"/>
	Min transmit power	18
	Max transmit power	Max
	Client aware	Enabled
	Scanning	Enabled
	Wide channel bands	5 GHz
	80MHz support	Disabled

Step 3: On the **Reboot Required** page, click **OK**.

Note Do not reboot the AP now, because you will do it in a subsequent step.



3.3 Configure the power settings

In this procedure, you configure the power settings for the 5-GHz and 2.4-GHz radios. In the 2.4-GHz band, set minimum power threshold to 6 and the maximum power to 9 for open-office and walled-office environments. In the 5-GHz band, for an open-office environment, set the minimum power threshold to 12 and the maximum to 15, or for a walled-office environment, set the minimum power threshold to 15 and the maximum to 18.

Enable background spectrum monitoring. When background spectrum monitoring is enabled, APs will continue to provide normal access service to clients. They will also monitor RF interference from neighboring APs and non-Wi-Fi sources, such as cordless phones and microwaves, on the channel they are servicing clients.

If you are in an open-office environment, choose option 1. If you are in a walled-office environment, choose option 2.

Option 1: Open-office environment

Step 1: Navigate to **RF > ARM**, click **Show advanced options**, and then click the **Radio** tab.

Step 2: In the **2.4 GHz band** section, click the **Customize ARM power range** checkbox, and then enter the following information.

- Background spectrum monitoring—**Enabled**
- Min power—**6**
- Max power—**9**

Step 3: In the **5 GHz band** section, click the **Customize ARM power range** checkbox, enter the following information, and then click **OK**.

- Background spectrum monitoring—**Enabled**
- Min power—**12**
- Max power—**15**

Note In all environments, the minimum power level differences between equal coverage level 2.4-GHz radios and 5-GHz radios should be 6 dBm.



The screenshot shows the 'RF' configuration window with the 'Radio' tab selected. The '5 GHz band' section is expanded, showing the following settings:

- Legacy only: Disabled
- 802.11d / 802.11h: Disabled
- Beacon interval: 100 ms
- Interference immunity level: 2
- Background spectrum monitoring: Enabled
- Customize ARM power range: ☒
- Min power: 12
- Max power: 15
- Smart antenna: Disabled
- ARM/WIDS Override: Dynamic
- Very high throughput: ☒

At the bottom of the window, there is a 'Hide advanced options' link and 'OK' and 'Cancel' buttons.

Option 2: Walled-office environment

Step 1: Navigate to **RF > ARM**, click **Show advanced options**, and then click the **Radio** tab.

Step 2: In the **2.4 GHz band** section, click the **Customize ARM power range** checkbox, and then enter the following information.

- Background spectrum monitoring—**Enabled**
- Min power—**6**
- Max power—**9**

Step 3: In the **5 GHz band** section, click the **Customize ARM power range** checkbox, enter the following information, and then click **OK**.

- Background spectrum monitoring—**Enabled**
- Min power—**15**
- Max power—**18**

The screenshot shows the 'RF' configuration window with the 'Radio' tab selected. The window is divided into two main sections: '2.4 GHz band' and '5 GHz band'. Each section contains a list of settings with dropdown menus and checkboxes. The 'Customize ARM power range' checkbox is checked in both sections. The 'Min power' and 'Max power' values are set to 6 and 9 for 2.4 GHz, and 15 and 18 for 5 GHz. The 'Background spectrum monitoring' is set to 'Enabled' in both sections. The 'Very high throughput' checkbox is checked in the 5 GHz section. At the bottom, there are buttons for 'OK' and 'Cancel', and a link for 'Hide advanced options'.

Setting	2.4 GHz band	5 GHz band
Legacy only:	Disabled	Disabled
802.11d / 802.11h:	Disabled	Disabled
Beacon interval:	100 ms	100 ms
Interference immunity level:	2	2
Background spectrum monitoring:	Enabled	Enabled
Customize ARM power range:	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Min power:	6	15
Max power:	9	18
Smart antenna:	Disabled	Disabled
ARM/WIDS Override:	Dynamic	Dynamic
Very high throughput:	<input type="checkbox"/>	<input checked="" type="checkbox"/>

3.4 Configure the employee SSID name and VLAN

This procedure configures the employee SSID name and VLAN for the IAP virtual controller. DMO is enabled to allow the AP to convert the multicast traffic to unicast for each client device. Unicast packets are transmitted at the higher unicast rate which decreases the airtime utilization and increases overall throughput. The default DMO channel utilization threshold is 90% which should be high enough to match the expected number of clients on the AP.

Step 1: In the **Network** section of the home page on the left side, under the **Name** column, click **New**.

Step 2: On the **WLAN Settings** page, click **Show advanced options**, enter the following information, and then click **Next**.

- Name—**Example-Employee**
- Primary Usage—**Employee**
- Dynamic multicast optimization—**Enabled**
- DMO channel utilization threshold—**90 %** (default)

WLAN Settings

Name & Usage

Name:

Primary usage: ☒ Employee ☐ Voice ☐ Guest

Broadcast/Multicast

Broadcast filtering:

Multicast transmission optimization:

Dynamic multicast optimization:

DMO channel utilization threshold: %

Transmit Rates

2.4 GHz: Min: Max:

5 GHz: Min: Max:

802.11

Band:

DTIM interval:

Min RSSI for probe request:

Min RSSI for auth request:

Bandwidth Limits

☐ Airtime ☐ Each radio

Downstream: kbps ☐ Per user

Upstream: kbps ☐ Per user

WMM

Background WMM: %

Best effort WMM: %

Video WMM: %

Voice WMM: %

Share

%

DSCP Mapping

Traffic Specification (TSPEC)

☐ Traffic Specification (TSPEC):

TSPEC Bandwidth: Kbps

☐ Spectralink Voice Protocol (SVP):

Miscellaneous

Content filtering:

Inactivity timeout: sec.

Deauth inactive clients:

SSID: ☐ Hide ☐ Disable

Out of service (OOS):

[Hide advanced options](#)

Step 3: On the **VLAN** page, enter the following information, and then click **Next**.

- Client IP assignment—**Network assigned**
- Client VLAN assignment—**Static**
- VLAN ID—**30**

The screenshot shows a configuration interface with four tabs: 1 WLAN Settings, 2 VLAN (active), 3 Security, and 4 Access. The main section is titled "Client IP & VLAN Assignment". It contains three settings: "Client IP assignment:" with radio buttons for "Virtual Controller managed" and "Network assigned" (selected); "Client VLAN assignment:" with radio buttons for "Default", "Static" (selected), and "Dynamic"; and "VLAN ID:" with a text box containing the value "30". At the bottom right, there are three buttons: "Back", "Next", and "Cancel".

3.5 Configure the employee SSID security

Employee security for the Wi-Fi network can be done with a WPA-2 personal passphrase or you can choose to have every employee authenticate with a username and password using WPA-2 enterprise. In either case, enable 802.11k and 802.11v for fast roaming. If you are planning to use WPA-2 enterprise, you should also enable opportunistic key caching to allow the APs to exchange pairwise master keys among themselves for faster roaming.

Note WPA-2 Enterprise is used to enable 802.1X authentication for a wireless network. The wireless client authenticates against the RADIUS server using an EAP-TLS exchange, and the IAP acts as a relay. Both the client and the RADIUS server use certificates to verify their identities.

With certain operating systems, the certificate is not automatically imported from the RADIUS server and requires manual installation in order for WPA-2 Enterprise to work. If the certificate is self-signed and generated on the RADIUS server, the certificate must be exported from the RADIUS server. From a Windows client, the certificate must be imported into the **Trusted Root Certification Authorities** store



If you are planning to use WPA-2 personal with passphrase access, choose option 1. If you are planning to use WPA-2 enterprise authentication, choose option 2.

Option 1: WPA-2 personal with passphrase access

Step 1: On the **Security** page, enter the following information, and then click **Next**.

- Security Level slider—**Personal**
- Key management—**WPA-2 Personal**
- Passphrase format—**8-63 chars** (default)
- Passphrase—**[password]**
- Retype—**[password]**
- 802.11k—**Selected**
- 802.11v—**Selected**

The screenshot shows the 'Security' tab in a configuration interface. At the top, there are four tabs: '1 WLAN Settings', '2 VLAN', '3 Security' (active), and '4 Access'. The 'Security Level' section on the left features a vertical slider with 'More Secure' at the top and 'Less Secure' at the bottom. The slider is positioned between 'Enterprise' and 'Open', with 'Personal' selected. The main configuration area on the right contains the following settings:

- Key management: WPA-2 Personal (dropdown)
- Passphrase format: 8-63 chars (dropdown)
- Passphrase: [masked with dots]
- Retype: [masked with dots]
- MAC authentication: Disabled (dropdown)
- Blacklisting: Disabled (dropdown)
- Enforce DHCP: Disabled (dropdown)
- Fast Roaming** section:
 - 802.11r: ☐
 - 802.11k: ☒
 - 802.11v: ☒

At the bottom right, there are three buttons: 'Back', 'Next', and 'Cancel'.

Option 2: WPA-2 enterprise with username and password

Step 1: On the **Security** page, enter the following information, and then click **Next**.

- Security Level slider—**Enterprise**
- Key management—**WPA-2 Enterprise**
- Authentication server 1—**Example-Radius**
- Opportunistic Key Caching (OKC)—**Selected**
- 802.11k—**Selected**
- 802.11v—**Selected**

The screenshot shows the 'Security' tab in the Aruba WLAN configuration interface. The 'Security Level' slider is set to 'Enterprise'. The 'Key management' dropdown is set to 'WPA-2 Enterprise'. The 'Authentication server 1' dropdown is set to 'Example-Radius'. The 'Authentication server 2' dropdown is set to '-- Select Server --'. The 'EAP offload' dropdown is set to 'Disabled'. The 'Reauth interval' is set to '0 min.'. The 'Authentication survivability' dropdown is set to 'Disabled'. The 'MAC authentication' section has two checkboxes: 'Perform MAC authentication before 802.1X' and 'MAC authentication fail-thru', both of which are unchecked. The 'Accounting' dropdown is set to 'Disabled'. The 'Blacklisting' dropdown is set to 'Disabled'. The 'Enforce DHCP' dropdown is set to 'Disabled'. The 'Fast Roaming' section has a checkbox for 'Opportunistic Key Caching (OKC)' which is checked. Below this, there are three checkboxes for '802.11r', '802.11k', and '802.11v', all of which are checked. At the bottom right, there are 'Back', 'Next', and 'Cancel' buttons.

Tab	Label
1	WLAN Settings
2	VLAN
3	Security
4	Access

Security Level

More Secure
Less Secure

Enterprise
Personal
Open

Key management: WPA-2 Enterprise

Authentication server 1: Example-Radius Edit

Authentication server 2: -- Select Server --

EAP offload: Disabled

Reauth interval: 0 min.

Authentication survivability: Disabled

MAC authentication:
☐ Perform MAC authentication before 802.1X
☐ MAC authentication fail-thru

Accounting: Disabled

Blacklisting: Disabled

Enforce DHCP: Disabled

Fast Roaming

Opportunistic Key Caching (OKC): ☒

802.11r: ☐

802.11k: ☒

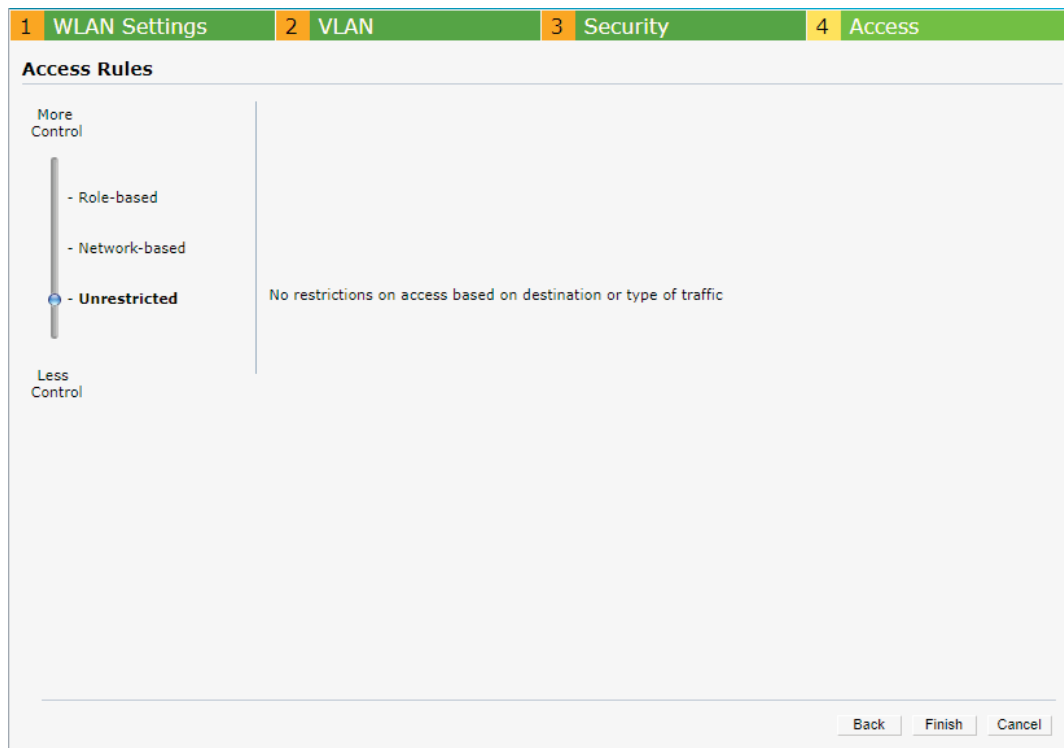
802.11v: ☒

Back Next Cancel

3.6 Configure the employee SSID access rules

This procedure configures the employee SSID access rules. In most cases, you will provide unrestricted access for your employees.

Step 1: On the **Access** page, move the **Access Rules** slider to **Unrestricted**, and then click **Finish**.



3.7 Configure the guest SSID name and VLAN

This procedure configures the guest SSID name and VLAN for the IAP virtual controller.

Step 1: In the **Network** section of the home page on the left side, under the **Name** column, click **New**.

Step 2: On the **WLAN Settings** page, enter the following information, and then click **Next**.

- Name—**Example-Guest**
- Primary usage—**Guest**

The screenshot shows a web interface for configuring WLAN settings. At the top, there are four tabs: '1 WLAN Settings' (highlighted in green), '2 VLAN', '3 Security', and '4 Access'. Below the tabs, the 'WLAN Settings' section is visible. It contains a 'Name & Usage' heading. Under 'Name:', there is a text input field containing 'Example-Guest'. Under 'Primary usage:', there are three radio button options: 'Employee', 'Voice', and 'Guest'. The 'Guest' option is selected. At the bottom left, there is a link 'Show advanced options'. At the bottom right, there are 'Next' and 'Cancel' buttons.

Step 3: On the **VLAN** page, enter the following information, and then click **Next**.

- Client IP assignment—**Network assigned**
- Client VLAN assignment—**Static**
- VLAN ID—**40**

The screenshot shows a configuration window with a tabbed interface at the top. The tabs are labeled '1 WLAN Settings', '2 VLAN', '3 Security', and '4 Access'. The '2 VLAN' tab is currently selected and highlighted in green. Below the tabs, the title 'Client IP & VLAN Assignment' is displayed. The configuration area contains three sections: 'Client IP assignment:' with radio buttons for 'Virtual Controller managed' and 'Network assigned' (the latter is selected); 'Client VLAN assignment:' with radio buttons for 'Default', 'Static' (selected), and 'Dynamic'; and 'VLAN ID:' with a text input field containing the value '40'. At the bottom right of the window, there are three buttons: 'Back', 'Next', and 'Cancel'.

3.8 Configure the guest SSID security

You can use WPA-2 personal encrypted passphrase for all your guests or you can require them to authenticate with a unique username and password. If you choose to require a passphrase, the most common captive portal is a simple acknowledgement splash page detailing the terms and conditions for using the guest network.

If you want to require a username and password, an open network is normally used to allow access to the captive portal. The authenticated splash page requires users to enter their personal details before being allowed onto the network. The username and passwords can be stored on the internal server in the virtual controller or in a RADIUS server.

Choose the guest SSID security option that you want to implement:

- If you are planning to use WPA-2 Personal encryption with passphrase access, choose option 1.
- If you are planning to use an open SSID with username and password authentication on a RADIUS server, choose option 2.
- If you are planning to use an open SSID with username and password authentication on the internal server, choose option 3.

Option 1: WPA-2 personal encryption with pre-shared key access

Step 1: On the **Security** page, enter the following information, and then click **Next**.

- Splash page type—**Internal - Acknowledged**
- Encryption—**Enabled**
- Key management—**WPA-2 Personal**
- Passphrase format—**8-63 chars** (default)
- Passphrase—**[password]**
- Retype—**[password]**

The screenshot shows the Aruba configuration interface with the **Security** tab selected. The interface is divided into two main sections: **Security Level** and **Splash Page Visuals**.

Security Level

- Splash page type: **Internal - Acknowledged**
- Captive portal proxy server:
- MAC authentication: **Disabled**
- Blacklisting: **Disabled**
- Enforce DHCP: **Disabled**
- Disable if uplink type is: ☐ 3G/4G ☐ Wifi ☐ Ethernet
- Encryption: **Enabled**
- Key management: **WPA-2 Personal**
- Passphrase format: **8-63 chars**
- Passphrase: *********
- Retype: *********

Splash Page Visuals

- Logo image: **No logo image** [Upload](#) | [Delete](#)
- Click thumbnail above to edit [Preview](#)
- Redirect URL: (optional)

At the bottom right, there are buttons for **Back**, **Next**, and **Cancel**.

Option 2: Open and authenticated with RADIUS server access

Step 1: On the **Security** page, enter the following information, and then click **Next**.

- Splash page type—**Internal - Authenticated**
- Auth server 1—**Example-Radius**

The screenshot shows the 'Security Level' configuration page, which is part of a four-step process (1. WLAN Settings, 2. VLAN, 3. Security, 4. Access). The page is divided into two main sections: 'Security Level' on the left and 'Splash Page Visuals' on the right.

Security Level Section:

- Splash page type:** Internal - Authenticated (dropdown menu)
- Captive portal proxy server:** (empty text field)
- WISPr:** Disabled (dropdown menu)
- MAC authentication:** Disabled (dropdown menu)
- Auth server 1:** Example-Radius (dropdown menu) with an **Edit** button next to it.
- Auth server 2:** -- Select Server -- (dropdown menu)
- Reauth interval:** 0 hrs. (text field and dropdown menu)
- Accounting:** Disabled (dropdown menu)
- Blacklisting:** Disabled (dropdown menu)
- Enforce DHCP:** Disabled (dropdown menu)
- Disable if uplink type is:** 3G/4G, Wifi, Ethernet (checkboxes)
- Encryption:** Disabled (dropdown menu)

Splash Page Visuals Section:

- Logo image:** No logo image (text) with **Upload** and **Delete** links.
- Click thumbnail above to edit** (text) with a **Preview** link.
- Redirect URL:** (text field) (optional)

At the bottom right, there are three buttons: **Back**, **Next**, and **Cancel**.

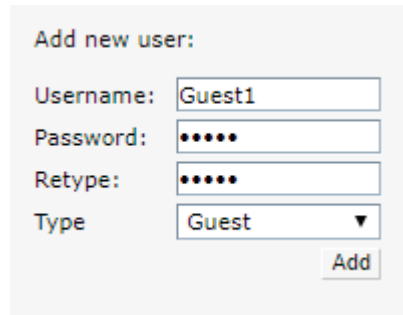
Option 3: Open and authenticated with internal server access

Step 1: On the **Security** page, enter the following information, and then click **Users**.

- Splash page type—**Internal - Authenticated**
- Auth server 1—**InternalServer**

Step 2: On the right side of the **Users** popup page, enter the following information, and then click **Add**.

- Username—**[Guest1]**
- Password—**[12345]**
- Retype—**[12345]**
- Type—**Guest**



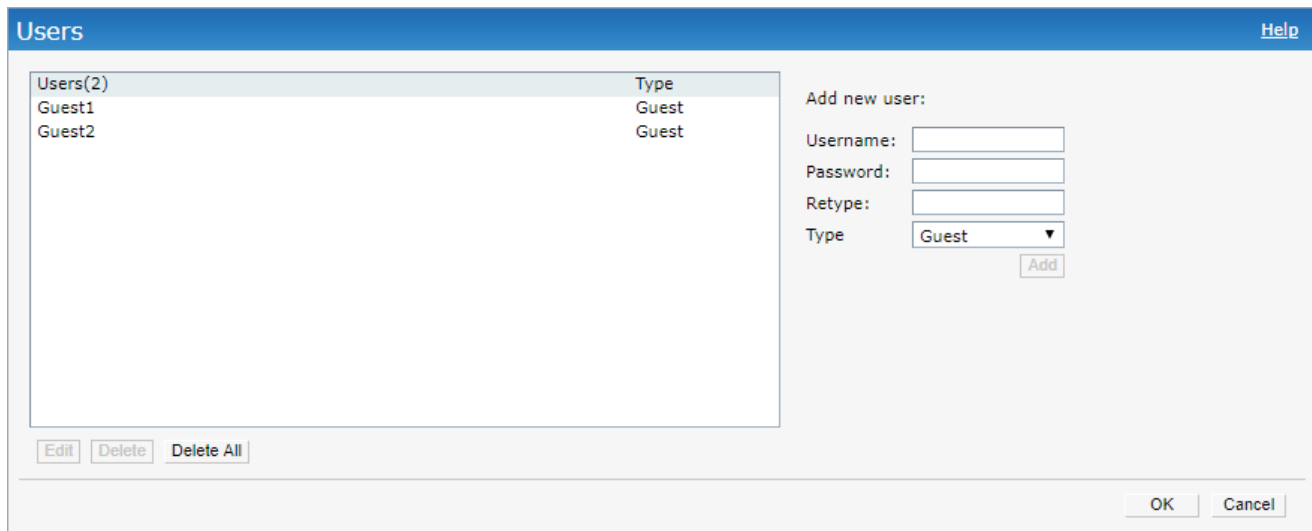
The screenshot shows a form titled "Add new user:" with the following fields and values:

Field	Value
Username:	Guest1
Password:	•••••
Retype:	•••••
Type	Guest

An "Add" button is located at the bottom right of the form.

Step 3: For each guest user, repeat Step 2.

Step 4: After adding the guest users, click **OK**.



The screenshot shows the "Users" popup page. On the left, there is a table with the following data:

Users(2)	Type
Guest1	Guest
Guest2	Guest

Below the table are buttons for "Edit", "Delete", and "Delete All". On the right side of the page is the "Add new user:" form, which is identical to the one in Step 2. At the bottom right of the popup are "OK" and "Cancel" buttons.

Step 5: On the **Security** page, click **Next**.

3.9 Configure the guest SSID access rules

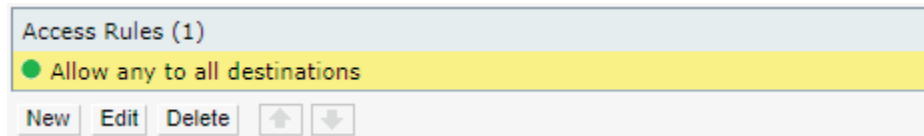
This procedure configures the guest SSID access rules. In most cases, you will provide access to DHCP and DNS services, and allow HTTP/HTTPS access to all destinations on the Internet.

The following table lists the access rules for guests.

Table 10 Access rules for guests

Rule Type	Service type	Service name	Action	Destination
Access control	Network	DHCP	Allow	To all destinations
Access control	Network	DNS	Allow	8.8.8.8 (well-known DNS)
Access control	Network	HTTP	Allow	To all destinations
Access control	Network	HTTPS	Allow	To all destinations
Access control	Network	Any	Deny	To all destinations

Step 1: On the **Access** page, move the **Access Rules** slider to **Network-based**, select the **Allow any to all destinations** rule, and then click **Delete**.



Step 2: In the **Access** page under **Access Rules**, click **New**, enter the following information, and then click **OK**.

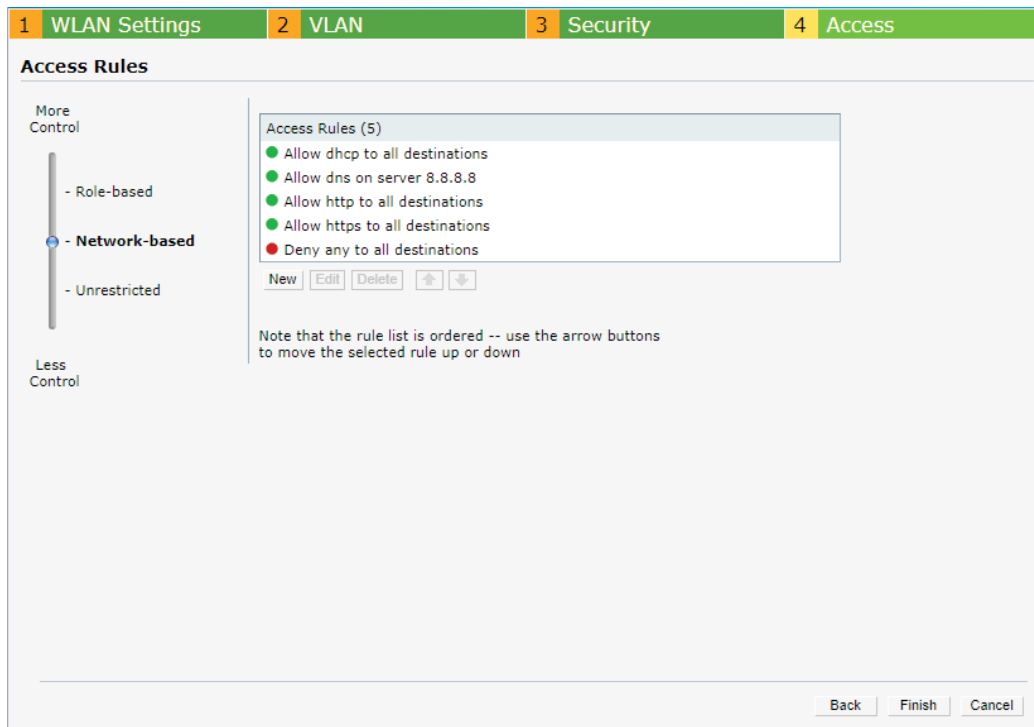
- Rule type—**Access control**
- Service type—**Network**
- Service name—**DHCP**
- Action—**Allow**
- Destination—**to all destinations**

A screenshot of the 'New Rule' dialog box. It contains several fields and options:

- Rule type:** A dropdown menu set to 'Access control'.
- Service:** A section with radio buttons. 'Network' is selected. Other options are 'Application', 'Application category', 'Web category', and 'Web reputation'.
- Service name:** A dropdown menu set to 'dhcp'.
- Action:** A dropdown menu set to 'Allow'.
- Destination:** A dropdown menu set to 'to all destinations'.
- Options:** A section with checkboxes for 'Log', 'Blacklist', 'Classify media', 'Disable scanning', 'DSCP tag', and '802.1p priority'. All are currently unchecked.
- At the bottom right are 'OK' and 'Cancel' buttons.

Step 3: For each rule in Table 10, repeat Step 2.

Step 4: In the **Access** page, click **Finish**.



3.10 Configure the access point name

Use this procedure to configure the access point name and repeat it for each AP. If you want this AP to be the virtual controller, enable the preferred master option.

Step 1: In the **Access Point** section of the home page, under the **Name** column select the name of the IAP, and then click **Edit**.

Step 2: On the **Edit Access Point** *Name* page under the **General** tab, enter the following information, and then click **OK**.

- Name—**305-AG1-AC1-21**
- [Optional] Preferred master—**Enabled**

The screenshot shows a dialog box titled "Edit Access Point 305-AG1-AC1-21" with a "Help" link in the top right corner. The "General" tab is active, showing the following fields and options:

- Name:** 305-AG1-AC1-21
- Zone:** (empty text box)
- Preferred master:** Enabled (dropdown menu)
- IP address for Access Point:**
 - ☒ Get IP address from DHCP server
 - ☐ Specify statically

At the bottom right, there are "OK" and "Cancel" buttons.

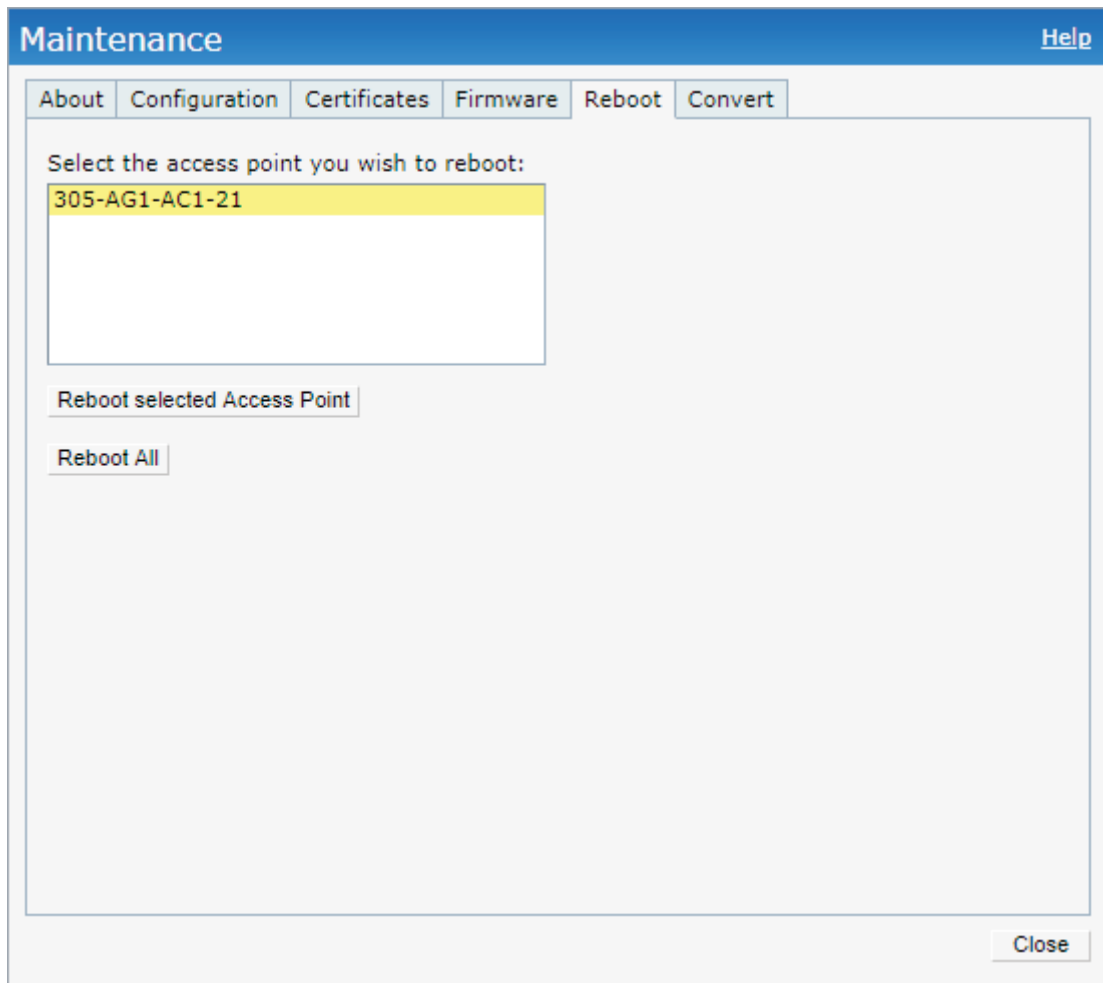
Step 3: On the **Reboot Required** message box, click **OK**.

The screenshot shows a "Reboot Required" message box with the following text:

Reboot Access Point a8:bd:27:c4:b5:14 for changes to take effect.

An "OK" button is located at the bottom right of the message box.

Step 4: Navigate to **Maintenance > Reboot**, highlight the new AP, and then click **Reboot selected Access Point**.



Step 5: For each IAP that you add to your network, repeat this procedure.

3.11 Remove the SetMeUp SSID

This procedure removes the initial SetMeUp SSID. If the SSID has not been automatically removed by the virtual controller, the procedure only has to be performed once.

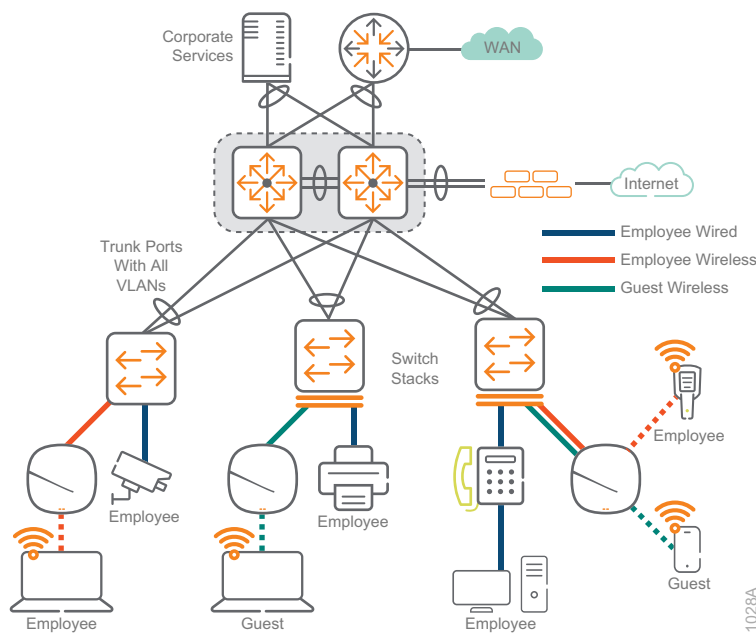
If the SetMeUp SSID has clients connected, it will remain active. To remove it manually, perform the following step.

Step 1: In the **Networks** section of the home page, select the **SetMeUp** SSID, and then click the **x**.

Summary

The flow of information is a critical component to a well-run organization. The Aruba Mobile First Campus design provides a prescriptive solution, based on best practices and tested topologies. This allows you to build a robust network that accommodates your organization's requirements. Whether users are located at a large LAN location or at a smaller remote site, this design provides a consistent set of features and functionality for network access, in order to help improve user satisfaction and productivity while reducing operational expense.

Figure 18 Mobile First Design



The Mobile First Campus design provides a consistent and scalable methodology of building your network, improving overall usable network bandwidth and resilience and making the network easier to deploy, maintain, and troubleshoot.

Validated Hardware and Software

The following list of hardware and software was validated for this guide:

Wired Aggregation

Product name	Software version
Aruba 5400R	16.04.0010
Aruba 3810M	16.04.0010

Wired Access

Product name	Software version
Aruba 5400R	16.04.0010
Aruba 3810M	16.04.0010
Aruba 2930M	16.04.0010
Aruba 32930F	16.04.0010

Wireless Access Points

Product name	Software version
Aruba 330 Series AP	6.5.4.5 (InstantOS)
Aruba 310 Series AP	6.5.4.5 (InstantOS)
Aruba 300 Series AP	6.5.4.5 (InstantOS)

What's New in This Version

The following changes have been made since Aruba last published this guide.

- This is a new guide



You can use the [feedback form](#) to send suggestions and comments about this guide.