

ARUBA CAMPUS FOR LARGE NETWORKS

Design & Deployment Guide

November 2019

Table of Contents

Document Conventions	1
Introduction	2
Purpose of This Guide.....	2
Customer Use Cases	3
Aruba Campus Design	5
Campus Wireless LAN Design Using Mobility Controllers	6
Wireless Design Components	22
Campus Wired LAN Design.....	26
Wired Design Components.....	40
Deploying the Aruba Campus	45
Campus Wired LAN.....	46
Configuring the ArubaOS-Switch Access Switch	47
Configuring the ArubaOS-Switch Aggregation Switch.....	63
Configuring the ArubaOS-CX Aggregation Switch.....	77
Configuring the Standalone Core Switches	87
Campus Wireless LAN	94
Configuring the Mobility Master	95
Configuring the Mobility Controller Cluster	105
Configuring the Wireless LANs.....	122
Configuring the Access Points.....	136
Summary	147
Validated Hardware and Software	148
What's New in This Version	149

Document Conventions

Bold text indicates a command, navigational path, or a user interface element. Examples:

- the **show stacking** command
- Navigate to **Configuration > System > General**
- click **Save**

Italic text indicates the definition of important terminology. Example:

- *Spatial streaming* is a transmission technique in MIMO wireless communication

Blue text indicates a variable for which you should substitute a value appropriate for your environment. Example:

- stacking member **2** priority **250**

Highlighting indicates emphasis. Example:

- ip address **10.4.20.2/22**

Note Notes contain asides or tips.



Caution Cautions warn you about circumstances that could cause a failure.



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 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 and switching products to create a high performance, secure and resilient campus network that is 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 security challenge for IT departments due to the lack of control over the devices. In addition, many IoT devices connect wirelessly to today's networks. IoT devices, such as building control systems, card readers, thermostats, and surveillance cameras, 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 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 a single system that integrates access points (APs), access switches, aggregation switches, core 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.

This 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 large organization with 500 to 10,000 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 and 802.11ax Wi-Fi supports speeds greater than 1 Gbps. To accommodate the increased data rates, the APs implement the 802.3bz standard of 2.5 and 5 Gbps. You can achieve the higher data rates on existing building twisted-pair cabling when connecting to Aruba switches with HPE Smart Rate ports. To support the explosion of IoT devices and latest wireless technologies, IEEE 802.3bt Power over Ethernet (PoE) provides simplicity and cost savings by eliminating the need for dedicated power. The access layer acts as a collection point for high-performance wired and wireless devices and must have enough capacity to support the power and bandwidth needs of today as well as scale for the future as the number of 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. In networks where 80% or more of the users are connecting via wireless but wired IoT devices continue to rise, the number of wired ports in the network is getting close to one per user.

This guide discusses 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
- High-performance core for non-stop forwarding of critical traffic
- IP multicast to efficiently propagate streaming traffic across the network

Aruba Campus Design

This design is targeted at large organizations supporting up to 10,000 users with multiple devices per user. The network could be a few floors in a building, a single building, or a group of 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. The core layer is used to interconnect aggregation-layer switches from multiple buildings or multiple floors in a building.

The three-tier design is used when there are several buildings in a campus that need to be connected and the number of aggregation switches or the layout of the physical wiring plant makes more sense to connect everything to a central core. For a network of 500 to 10,000 users, the three-tier campus design is the most common, as shown in the following figure.

Figure 1 Three-tier campus design



The Aruba Campus design uses access switches or switch stacks connected to a dual-switch aggregation layer. 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 are less expensive than a stack of fixed access switches. The aggregation layer is dual-connected into a pair of high-speed core switches, which provide a maximum level of redundancy and resiliency for non-stop forwarding.

The aggregation layer also provides critical network services like WAN aggregation, Internet DMZ and data center servers for an organization. The aggregation switches need numerous 10 Gbps ports into the access layer and 40 Gbps ports into the core. The high-speed core switches must support a large number of 40 Gbps ports with enough capacity to grow and several 100 Gbps ports to connect between themselves.

The Aruba Campus design uses Aruba APs and Mobility Controllers for wireless access because they provide ease of configuration and maximum operational flexibility. This design minimizes the number of different components in order to make operations, maintenance, and troubleshooting simpler.

CAMPUS WIRELESS LAN DESIGN USING MOBILITY CONTROLLERS

The Aruba 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 Aruba wireless solution include the following:

- Location-independent network access improves employee experience and productivity.
- Hard-to-wire locations receive network connectivity without costly construction.
- Wireless is plug-and-play; the network automatically recognizes and provisions new APs.
- Reliable and high-performance wireless connectivity, including automated radio frequency (RF) spectrum management.
- Application visibility and control can identify and prioritize business critical applications.
- Centralized control of wireless environment allows easy management and operation.
- Pay as you grow with controller clustering for increasing network capacity and seamless failover.
- Live upgrades perform operating system updates, and in-service module updates support 24/7 operations.

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 consistent performance and better connection reliability.

Aruba Mobility Master and Mobility Controllers

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*. The Aruba Mobility Master and mobility controllers offer centralized network engineering, IP services, security, and app-aware policy controls. The mobility controller is responsible for many of the operations that traditionally would be handled by an autonomous AP, and it delivers additional functionality for control, security, operation, and troubleshooting.

Mobility Master Features

The Aruba WLAN design for large campus incorporates a centralized Mobility Master, which is a next-generation controller that provides simplified operations and enhanced performance by centrally managing mobility controllers. The Mobility Master provides the following advanced features:

- **High scale and reliability**—Mobility Master can centrally manage up to 100,000 clients, 10,000 APs and 1,000 mobility controllers. Mobility controllers become managed devices under the Mobility Master and handle the AP terminations.
- **24/7 always-on operations**—Live upgrades and multiple software version support eliminate the need for planned maintenance windows or downtime. You can selectively upgrade each controller cluster or individual service modules—such as AppRF, AirGroup, AirMatch, north-bound APIs, Unified Communications Manager, and web content classification that resides on the Mobility Master—without requiring an entire system reboot.
- **Automated RF management**—AirMatch is an enhancement to Adaptive Radio Management technology. AirMatch ensures even channel use, assists in interference mitigation, and maximizes system capacity by automating channel selection and channel-widths and optimizing transmit power.
- **Flexibility of deployment**—You can deploy Mobility Master by using a virtual machine (VM) or an x86-based hardware appliance. If you already have a VM environment, you can benefit from ease of operation and right-size your VM by adjusting the CPU or memory. Moving to a VM-based deployment that has more memory and compute resources allows you to manage more services on the network. The virtual Mobility Master can run on open source KVM, VMWare ESXi, or Microsoft HyperV hypervisors.

Mobility Master Configuration Hierarchy

The Mobility Master is a centralized management platform in a multi-tier architecture that provides separation of management, control, and forwarding plane. It maintains all configurations, including its own, eliminating multiple points of contact to apply global and local configurations to each managed device. You organize your common configurations at a higher level of the hierarchy, and they are propagated to the lower levels. You configure group-specific or device-specific changes at the lower levels if they do not pertain to the higher-level devices. This type of hierarchy simplifies the configuration of design elements that are shared across common deployment types.

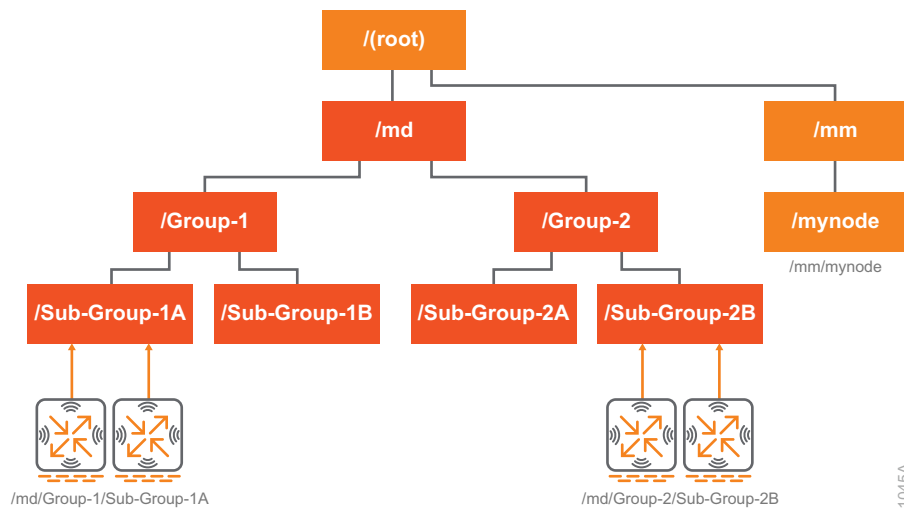
The following table shows the group structures for the hierarchy.

Table 1 Hierarchy structure

Category	Name	Description
Mobility Master	/	Configurations common to Mobility Master and its managed devices. NOTE: Configuration changes are not allowed at the root.
	/md	Configurations common to all managed devices. The user can create additional groups under this group.
	/md/<group name>	The group name is used to differentiate the devices physically or by the type of deployment, such as DMZ, Branch, Campus, RAPs, and so on.
	/mm	Configurations common to the primary and standby Mobility Master (VRRP pair).
	/mm/mynode	Configurations specific to a Mobility Master. This can only be edited on the respective Mobility Master.

The Mobility Master hierarchy streamlines the configuration process by supporting multiple configurations for multiple deployments at different locations. Configuration elements can be mapped to one or more end devices, such as a managed device or VPN concentrator. Common configurations across devices are extracted to a shared template, which merges with device-specific configurations to generate the configuration for an individual device. The network can be organized in a hierarchy of up to five levels, including groups, sub-groups, and the managed devices that are added to these groups. The following is an example of a configuration hierarchy with three levels.

Figure 2 Mobility Master configuration hierarchy



Mobility Controller Features

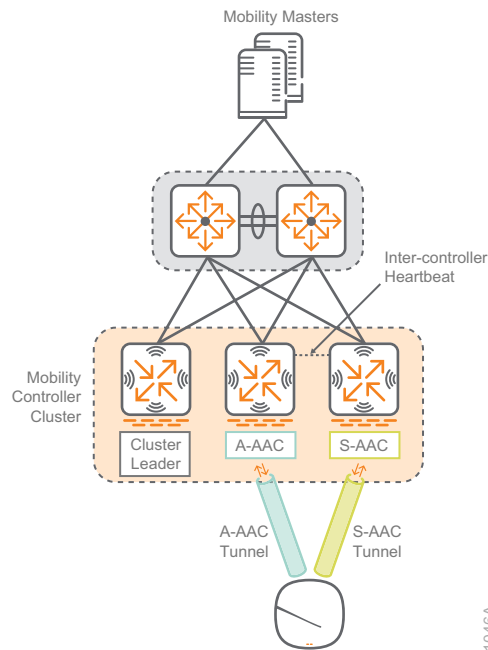
Mobility controllers are managed devices under the Mobility Master and handle the AP terminations. Clustering the mobility controllers provides high availability, scalability, and seamless roaming throughout an entire building or across many buildings in a campus. With clustering, users are not affected by a controller failure, and their traffic continues to flow without noticeable impact. The user session information is shared across controllers in the cluster to ensure there is no single point of failure for any user. Mobility controllers and clustering provide the following benefits:

- **Seamless roaming**—WLAN clients remain anchored to a single member of the cluster regardless of where they roam or to which AP they connect. Users can roam across entire campus without loss of firewall state or the need to re-authenticate or change their IP address.
- **Hitless failover and automated load balancing**—Users sessions and AP traffic are load-balanced across a cluster of mobility controllers to optimize network utilization during peak hours and preserve user experience during unplanned outages. In the event of a cluster member failure, connected clients fail over to a redundant cluster member. If needed, clients are moved among cluster members in a stateful manner to prevent congestion on a single controller and disruption of service.
- **Client performance optimization**—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 do not trigger AP changes for clients already associated with an AP.
- **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.
- **Intelligent Application Identification**—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. 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.

Mobility Controller Operations

After a cluster is formed, a leader is elected to perform several responsibilities, like managing AP and user anchor controller roles. An anchor controller is a single mobility controller that manages the session activity between individual APs and users. When redundancy is enabled, the cluster leader assigns an active AP anchor controller (A-AAC) and standby AAC (S-AAC) for each AP. The cluster leader also dynamically load-balances the APs among the active cluster members. Once the AAC roles are chosen, tunnels are formed between the AP and the active and standby AACs as depicted in the following figure.

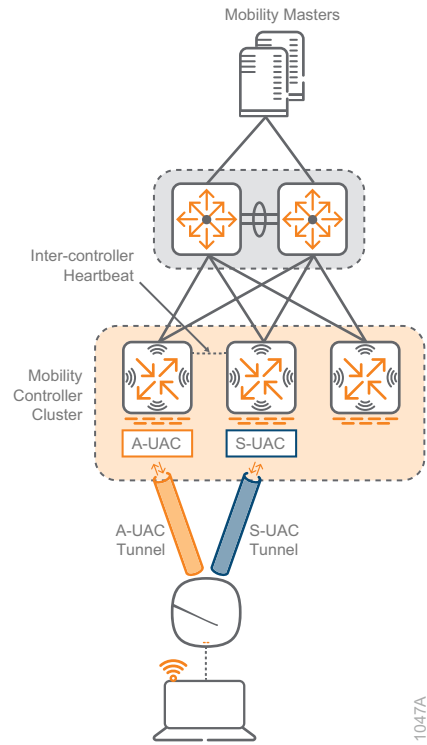
Figure 3 Cluster redundancy—active and standby AP anchor controllers



An inter-controller heartbeat keeps the two cluster controllers synchronized. If the AP's A-AAC fails or is taken out of service, the AP will failover to the S-AAC and the cluster leader will choose a new S-AAC.

The cluster leader also assigns an active user anchor controller (A-UAC) and standby UAC (S-UAC) for each WLAN user. If redundancy is disabled, the client stateful failover feature is not available. When a user joins the wireless network, tunnels are formed to the active and standby UACs as depicted in the following figure.

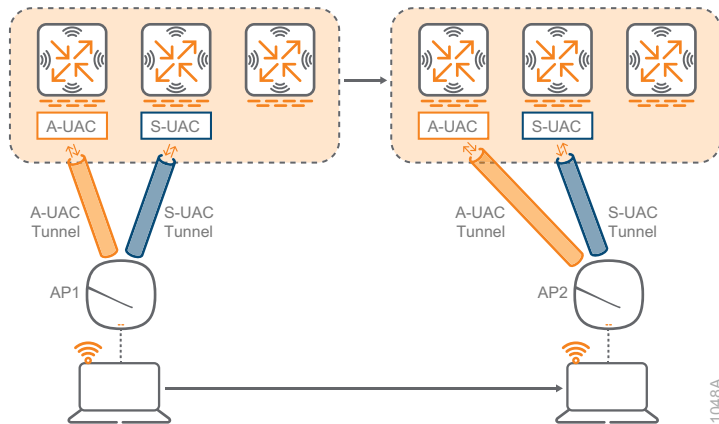
Figure 4 Cluster redundancy—active and standby user anchor controllers



An inter-controller heartbeat keeps the two cluster controllers synchronized. If the user A-UAC fails or is taken out of service, the user session is moved to the S-UAC and the cluster leader chooses a new S-UAC.

If the user roams to another AP, the tunnel is maintained to the original A-UAC and S-UAC even if the new AP has a different set of AACs. This means a user can seamlessly roam throughout an entire building or campus as depicted in the following figure.

Figure 5 Seamless roaming—active and standby user anchor controllers



The functionality provided by the mobility controller in this design includes:

- Acting as a user-based application firewall
- Terminating user-encrypted sessions from wireless devices
- Providing certificate-based IPsec security to protect control channel information
- Performing user authentication such as 802.1X and captive portal authentication
- Providing guest access and captive portal services
- Providing self-contained management by way of a Mobility Master hierarchy with the master pushing configuration to other mobility controllers to reduce administrative overhead
- Delivering AP software updates automatically when the mobility controller is upgraded

This level of seamless, integrated functionality eliminates many of the challenges experienced with traditional systems integration of these services. Network administrators need to learn only one interface, which reduces deployment complexity and speeds problem resolution across a broad range of designs and solutions.

Access Point Placement

Aruba recommends doing a site survey for all wireless network installations. The main goal of a site survey is to determine the feasibility of building a wireless network on your site. You also use the site survey to determine the best place for access points and other equipment, such as antennas and cables. With that in mind, the following guidelines can be used as a good starting point for most office environments.


For typical wireless bandwidth 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, leads to 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 2500-user network with three devices per person, this works out to 7500 total devices, and with 30 devices per AP, this translates to approximately 250 APs for this example.

Whenever possible, APs should be placed near users and devices in offices, meeting rooms, and common areas, instead of in hallways or closets. 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 6 Sample office AP layout (not to scale)



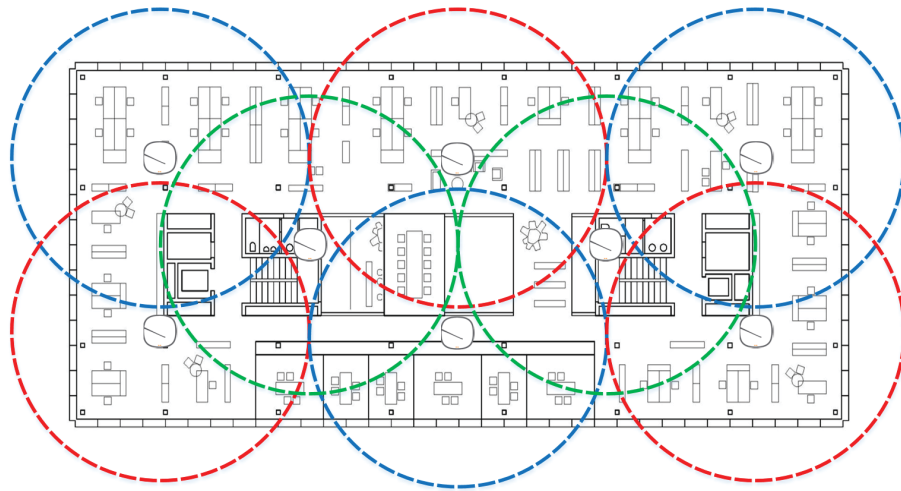
After studying your environment with the 35-50-foot (10-15 meter) rule in mind, make sure you also have enough capacity for the number of users. In an average office environment with APs every 35-50 feet (10-15 meters), the 30 devices per AP average will easily be satisfied. However, if you have high-density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

Channel Planning

The centralized Aruba AirMatch software is very good at automating channel assignment, and for most wireless installations, channel selection and transmit power can be left to its advanced 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 caused by two radios being on the same channel. We recommend using only these three channels for your 2.4-GHz installations in order to avoid the more serious problem of adjacent channel interference caused by radios on overlapping channels or adjacent channels with radios too close together. A professional site survey could further optimize this type of design with a custom power level, channel selection, and enabling and disabling 2.4 GHz radios for optimal coverage and to minimize interference.

Figure 7 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 8 802.11ac 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 DFS 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 stops transmitting on that channel for a time and moves 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. 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 vary from location to location.

You can find a list of the 5-GHz channels available in different countries 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))

Power Settings

The optimum power settings vary based on your physical environment, and you should always follow the recommendations from your professional site survey to finalize your initial settings. With that in mind, use the following guidelines for a typical wireless design.

- In the 2.4-GHz band, set the minimum power threshold to 6 dBm and the maximum power to 9 dBm for open-office and walled-office environments
- In the 5-GHz band, set the minimum power threshold to 12 dBm and the maximum to 15 dBm for an open-office environment
- In the 5-GHz band, set the minimum power threshold to 15 dBm and the maximum to 18 dBm for a walled-office environment
- In all environments, do not exceed a power level difference of 6 dBm between the minimum and maximum settings on all radio bands
- In all environments, the minimum power level differences between equal coverage level 2.4-GHz radios and 5-GHz radios should be 6 dBm

After the initial settings are configured, the AirMatch software automatically adjusts power settings as needed and on an ongoing basis.

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 2 Theoretical bandwidth for 802.11ac at various channel widths and spatial stream counts

Channel width	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	N/A	N/A

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.

Site Survey

A site survey is an important tool that gives you a solid understanding of the radio frequency behavior at your site and, more importantly, where and how much interference you might encounter with your intended coverage zones. A site survey also helps you to determine what type of network equipment you need, where it goes, and how it needs to be installed. A good survey allows identification of AP mounting locations, existing cable plants, and yields a plan to get the wireless coverage your network requires. RF interacts with the physical world around it, and because all office environments are unique, each wireless network has slightly different characteristics. The recommendations listed in the sections above are a good starting point, but a solid site survey allows you to customize the RF plan for your specific location.

If you want to provide ubiquitous multimedia coverage in a multi-floor/multi-building campus with uninterrupted service, you need a professional site survey to balance the elements required for success. Planning tools have evolved with the radio technologies and applications in use today, but a familiarity with the RF design elements and mobile applications is required to produce a good plan. Completing a site survey now yields good information that can be used again and again as the wireless network grows and continues to evolve.

RF Performance Optimization

After a successful site survey helps you properly place your APs, there are additional ways to provide long-term performance management for your wireless network. The Aruba AirMatch feature models the network and improves the user experience by optimizing the RF performance on an ongoing basis. The APs collect information about their RF neighborhood and forward the data to the AirMatch process running on the Mobility Master. AirMatch consumes the RF information from the entire network and generates a new RF plan each day. If the new plan has enough variance from the existing plan to warrant a change, an update is sent out to all APs at a pre-determined time. However, AirMatch can only optimize the RF environment to a certain degree if the APs are not initially located correctly in your environment.

Aruba also provides an emerging tool called NetInsight, which uses machine learning-based network analytics to deliver recommendations for optimization around mobile workers, wireless and IoT devices. Data from multiple sources including your wireless infrastructure, DHCP and authentication servers are gathered in an onsite data collector. The data is compressed and sent via a secure tunnel to the NetInsight cloud instance where network connectivity and performance issues are analyzed by leveraging machine learning-based models using Aruba's Wi-Fi expertise and the latest cloud technologies. A web-based dashboard allows you to view insights along with root causes, and more importantly, it provides recommendations to fix immediate and foreseeable network performance issues. Aruba 5xx series access points work seamlessly with NetInsight to automatically power down when connectivity demand ceases and power up when demand returns. NetInsight uses predictive analytics and machine learning to identify usage patterns. After a brief learning period, NetInsight can predict when demand stops and when it starts. Each Green AP-enabled access point can cut electricity usage from approximately 21 watts at full PoE power to just 6 watts in sleep mode, resulting in a savings of more than seventy percent.

Channel Planning 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 get better performance split between two radios on 20-MHz channels than they do on one radio using a 40-MHz channel. This is because the more channels you have to use, the better overall throughput is for a higher number of devices. As mentioned previously, a typical Aruba wireless installation uses the AirMatch software running on the Mobility Master and NetInsight running in the cloud for RF channel planning and optimization.

802.11ax (Wi-Fi 6) Enhancements

Designed to address connectivity issues for high density deployments, the new 802.11ax standard improves the performance of the entire network. New features allow multiple clients to transmit simultaneously, increasing network capacity by up to 4 times compared to 802.11ac.


The most significant new feature of the 802.11 ax standard is orthogonal frequency-division multiple access (OFDMA), which replaces orthogonal frequency-division multiplexing (OFDM). Other important new features include BSS coloring and the ability to transmit up to 8 clients with Multi-User Multiple Input Multiple Output (MU-MIMO).

OFDMA

With OFDM, frames are transmitted consecutively using the entire channel to a single client at a time. For example, if a client is connected to a 20 MHz wide channel and sends data, the entire channel is taken up, and then the AP and clients take turns, one at a time, sending data on the channel.

OFDMA changes that behavior. Channels are divided into smaller sub-channels, and the AP can send data to multiple clients at a time. A 20 MHz wide channel can support up to nine clients, and the number of sub-channels continually adjusts in order to support fewer higher-speed clients or more lower-speed clients. Sub-channel use is dynamic and adjusts automatically every transmission cycle, depending on client data needs.

Figure 9 OFDMA operation in 802.11 ax—multiple clients share the channel




Wider channels can support even more sub-channels. An 80-MHz-wide channel can support up to 37 clients at a time. OFDMA supports downlink traffic, from the AP to the clients, and will eventually support uplink traffic, from the clients to the AP.

8X MU-MIMO

The 802.11ax standard enhances MU-MIMO and supports up to eight clients at a time (the 802.11ac standard allowed for eight, but vendors implemented only four or less). This feature effectively doubles the number of devices to which an AP can talk.

Figure 10 8x8:8 MU-MIMO to single and dual stream clients



BSS Coloring

BSS coloring allows the network to assign a “color” tag to a channel and reduce the threshold for interference. Network performance improves because APs on the same channel can be closer together and still transmit at the same time if they are different colors. The field is 6-bits, so there are 63 different colors available.

Figure 11 BSS coloring—same channel only blocked on color match



Security

WPA3—Aruba Simultaneous Authentication of Equals (SAE) protocol was added in the late 2000s to the IEEE 802.11s (mesh networking) standard. IEEE 802.11s was certified in 2012. SAE is an instantiation of the dragonfly key exchange, which performs a password-authenticated key exchange by using a zero-knowledge proof—each side proves it knows the password without exposing the password or any password-derived data.

WPA3 introduces a new configuration option for 802.1X/EAP called Commercial National Security Algorithms (CNSA). CNSA was defined by the United States National Security Agency to protect secret and top-secret data on government and military networks. Since CNSA affords consistent security without the ability to misconfigure, it is being adopted by enterprises that have strong security requirements, such as financial institutions. CNSA establishes a suite of cryptographic algorithms that all afford roughly the same level of protection: SHA384 for hashing, NIST's p384 elliptic curve for key establishment and digital signatures, and AES-GCM-256 for data encryption and authentication. With CNSA, the EAP method must be EAP-TLS and the negotiated TLS cipher suite must exclusively use cryptographic algorithms from the CNSA suite.

Enhanced Open—Aruba Opportunistic Wireless Encryption (OWE) is an alternative to open networks. It has the same workflow and the same user requirements. Basically, click the available network and get connected. To the user, an OWE network looks just like an open network (with no padlock symbol), but the advantage is that it's encrypted. OWE performs an unauthenticated Diffie-Hellman key exchange when the client associates with the AP. The result of that exchange is a key known only to two entities in the entire world, the client and the AP. That key can be used to derive keys to encrypt all management and data traffic sent and received by the client and AP.

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.

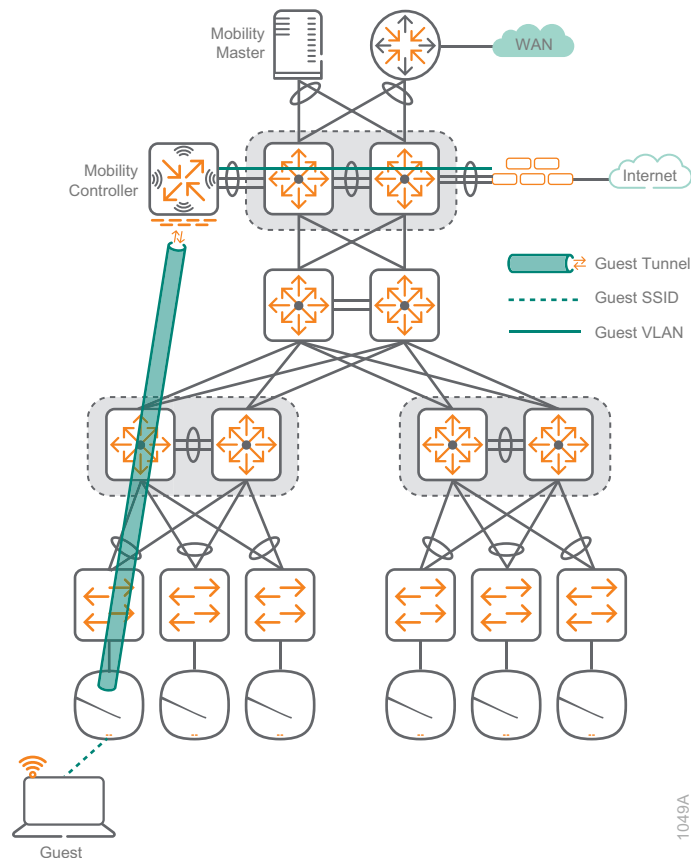
The flexibility of the Aruba Campus architecture allows the wireless network to provide guest and employee access over the same infrastructure. This integrated ability simplifies network operations and reduces capital and operational costs. The critical part of the architecture is to ensure that guest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

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

- Provides Internet access to guests through an open wireless Service Set Identifier (SSID), with web access control in the firewall.
- 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.

Every AP can be provisioned with controlled, open access to wireless connectivity and the Internet. From the wireless AP, guest traffic is securely tunneled back to the controller and 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 12 Guest wireless network



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 mobility 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.

Campus Wireless LAN Design Summary

The Aruba 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 Aruba wireless solution include:

- Seamless network access for employees, guests and IoT devices.
- Plug and play deployment for wireless APs.
- AirMatch and ClientMatch technology to maximize WLAN performance by dynamically choosing the best Wi-Fi channel, channel width, transmit power and client to radio steering.
- AppRF and Intelligent Application Identification to provide visibility into the applications running on the wireless network.
- Non-stop networking with a centralized Mobility Master and controller clustering for pay-as-you-grow and seamless failover.
- Live upgrades to perform operating system updates without an outage or service impact.
- In-service module upgrades to dynamically update individual service modules without requiring an entire system reboot.

WIRELESS DESIGN COMPONENTS

You can deploy Aruba wireless in two main modes, controller-based or controllerless. With Aruba's Mobility Master and mobility controllers, certain features run on the master while others run on the cluster of controllers. This type of design is typically used in larger networks with more than 500 users, but there is no reason it cannot be used in smaller networks.

Access Points

There are currently two series of Aruba access points: the latest generation 5xx series 802.11ax APs and the 3xx series 802.11ac Wave 2 APs. Details about currently available models are listed below; they support different throughput and client loads to meet different deployment needs.

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 current Aruba 5xx and 3xx APs:

- Unified AP for either controller-based or controllerless deployment modes
- Hitless PoE failover between both Ethernet ports (dual Ethernet models only)
- Built-in Bluetooth Low-Energy radio
- Advanced Cellular Coexistence to minimize interference from cellular networks
- Application visibility for QoS and traffic control
- Enhanced security with WPA3 and Enhanced Open

Aruba 5xx Series Access Point Options

The Aruba 5xx Series of campus access points support 802.11ax to efficiently and simultaneously serve multiple clients and traffic types in dense environments. These APs offer increased data rates for both individual device and overall system while delivering high performance and throughput in environments where mobile and IoT density is a growing concern.

Aruba 5xx common capabilities:

- Dual uplink ports with LACP support for redundancy and increased capacity
- Bluetooth 5 and Zigbee radios for location and IoT use-cases
- Green AP mode for energy savings up to 70%

Aruba 550 Series Access Points—The Aruba 550 Series APs are ideal for extreme high-density environments, such as public venues, higher education, hotels, and enterprise offices. The 550 series supports maximum data rates of 4.8Gbps in the 5GHz band and 1,150Mbps in the 2.4GHz band (for an aggregate peak rate of 5.95Gbps). The Aruba 550 series requires ArubaOS and Aruba InstantOS 8.5 software.

- Dual-radio (8x8 + 4x4 MIMO)
- Optional tri-radio mode* with two 5GHz and one 2.4GHz radio (all 4x4 MIMO)
- Dual 5G HPE Smart Rate ports
- AI-powered features for wireless RF and client connectivity optimization
- Up to 1024 associated client devices per radio (recommended active 200) *

**Some 5xx features are not supported in the initial release but will be enabled in future software releases.*

Aruba 530 Series Access Points—The Aruba 530 Series APs are ideal for very high-density environments, such as higher education, K12, retail branches, hotels and digital work places. The 530 series supports maximum data rates of 2.4Gbps in the 5GHz band and 1,150Mbps in the 2.4GHz band (for an aggregate peak rate of 3.55Gbps). The Aruba 530 series requires ArubaOS and Aruba InstantOS 8.5 software.

- Dual-radio (dual 4x4 MIMO)
- Dual 5G HPE Smart Rate ports
- AI-powered features for wireless RF and client connectivity optimization
- Up to 1024 associated client devices per radio (recommended active 200)*

**Some 5xx features are not supported in the initial release but will be enabled in future software releases.*

Aruba 510 Series Access Points—The Aruba 510 Series APs are ideal for high-density environments, such as schools, retail branches, hotels, and enterprise offices. The 510 series supports maximum data rates of 2.4Gbps in the 5GHz band and 575Mbps in the 2.4GHz band (for an aggregate peak data rate of 2.975Gbps). The Aruba 510 series requires ArubaOS and Aruba InstantOS 8.4 software.

- Dual-radio (4x4 + 2x2 MIMO)
- Single 2.5G HPE Smart Rate and Gigabit Ethernet uplink ports
- Up to 256 associated client devices per radio

Access 3xx Series Point Options

Aruba 340 Series Access Points—The Aruba 340 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
- 5G HPE Smart Rate and Gigabit Ethernet uplink ports with Link Aggregation Control Protocol (LACP) support for increased capacity
- Hitless 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
- 2.5G HPE Smart Rate and Gigabit Ethernet uplink ports with LACP support for increased capacity

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.

Mobility Master and Mobility Controllers

You can deploy the Aruba Mobility Master either as a virtual appliance or on an x86-based hardware appliance. The Aruba 7200 series mobility controller is optimized for mobile application delivery in order to ensure the best mobility experience over Wi-Fi. The Aruba 7000 Series mobility controllers optimize cloud services and secure enterprise applications for hybrid WAN at branch offices, while reducing the cost and complexity of deploying and managing the network. The virtual mobility controllers run on x86 platforms, and you can integrate them into your existing virtual machine environments.

Mobility Master and Mobility Controller Options

The Mobility Master provides centralized configuration and visibility, multi-tenant with multizone, controller clustering with hitless failover, automatic user and AP load balancing and seamless roaming. They support live upgrades, multiple OS and high performance WiFi with AirMatch.

The mobility controllers support Aruba's Next-Generation Mobility Firewall with AppRF technology as well as other enterprise-critical capabilities like authentication, encryption, IPv4 and IPv6 services, Adaptive Radio Management, ClientMatch, and RFProtect spectrum analysis and wireless intrusion protection.

Table 3 Maximum scaling capabilities

Device	Mobility controllers	Access points	Clients	Cluster members
Mobility Master	1000	10,000	100,000	2
7200 Mobility Controllers	N/A	2048	32,768	12
7000 Mobility Controllers	N/A	64	4096	4
Mobility Controllers Virtual Appliance (MC-VA)	N/A	1000	16,000	4

CAMPUS WIRED LAN DESIGN

The campus LAN not only provides wired and wireless connectivity for local users but becomes the core for interconnecting 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. A consistent, well-structured modular design provides the highest level of availability, ease of deployment, and operations.

To accommodate growth in the number of devices, network engineers build wired LANs in layers. A typical wired LAN with 500 to 10,000 users has an access layer, an aggregation layer and a core layer. With the Aruba 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
- Consistent performance and user experience across the entire campus
- Minimize the impact of network failures by reducing the size of the fault domains

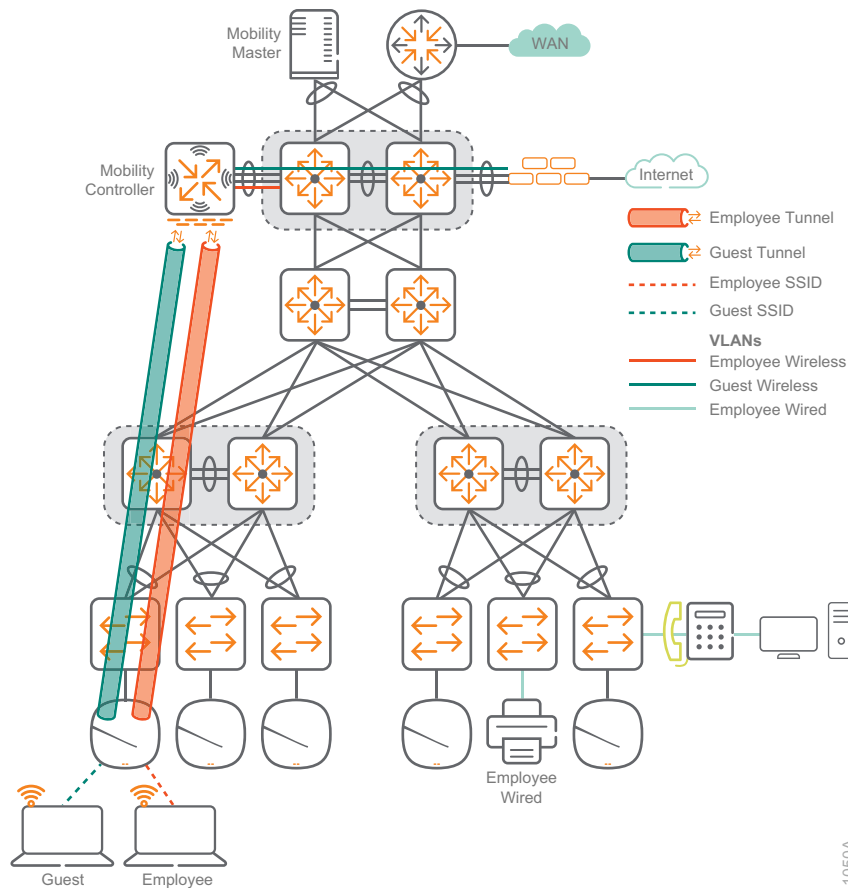
Access Layer

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. To simplify the network as much as possible, the Aruba Campus access layer is a layer-2 design. Aruba aggregation switches can easily accommodate thousands of devices, so the old rule of limiting VLANs 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. In this design, we use separate VLANs for employee wired, employee wireless, and guest wireless traffic. Employee wired traffic VLAN is used by the trusted devices cabled to the access layer switches. Employee wireless VLAN is used by trusted devices on Wi-Fi; notice in the following figure that the VLAN extends to the mobility controller over the Employee Tunnel. 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. Employee and guest traffic are segmented on the network, as depicted in the following figure.

Figure 13 Employee and guest VLANs



Access Layer Switching Features

The following features highlight several of the key aspects of the Aruba access layer design.


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 switch, 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. 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.

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 following figure. 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 14 Three-switch ring topology and roles



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. Be careful when deploying this feature when multiple MAC addresses are required behind a single port (for example, VMs, IP phones, unmanaged switches, etc.).
- **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.
- **Loop Protection**—Provides protection against loops by transmitting and monitoring of loop protocol packets. If a loop is detected the port that transmitted the loop protocol packet is disabled. This mechanism helps prevent loops if an unmanaged switch that does not support STP is connected to the network.
- **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 (RAs) and filters out the ones sent by unauthorized devices.

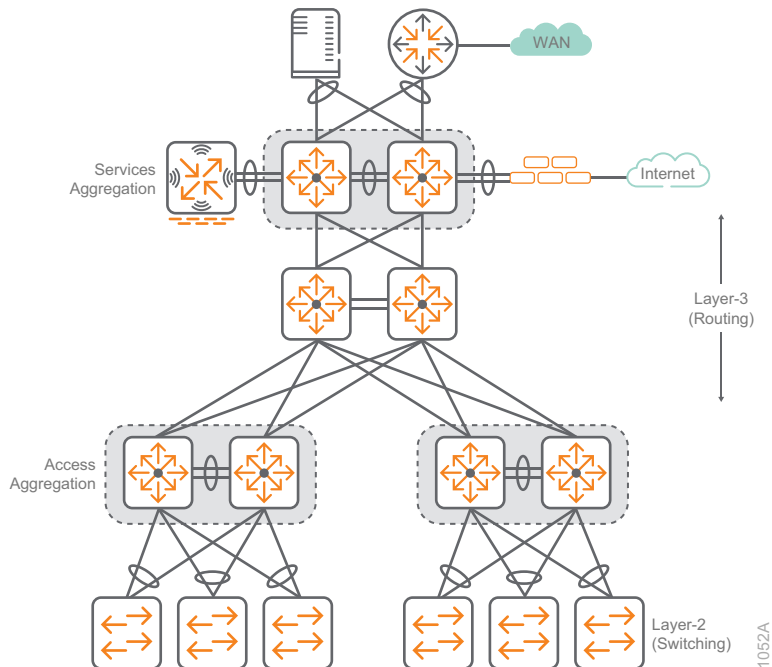
IP Multicast

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 services aggregation block also becomes the ideal location for connecting other network services, such as the WAN aggregation, Internet DMZ, and server rooms or data center for a large organization.

Figure 15 Aggregation layer—routing and switching boundary




Aggregation Layer Switching Features

The following features highlight key aspects of the Aruba aggregation layer design.

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 adjacent switches 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 blocks 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 300 ms.

Figure 16 Traditional spanning tree vs MC-LAG



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, virtual switching framework (VSF), or Virtual Switching Extension (VSX) in order to appear as a single switch to other layer-2 attached devices in the network.

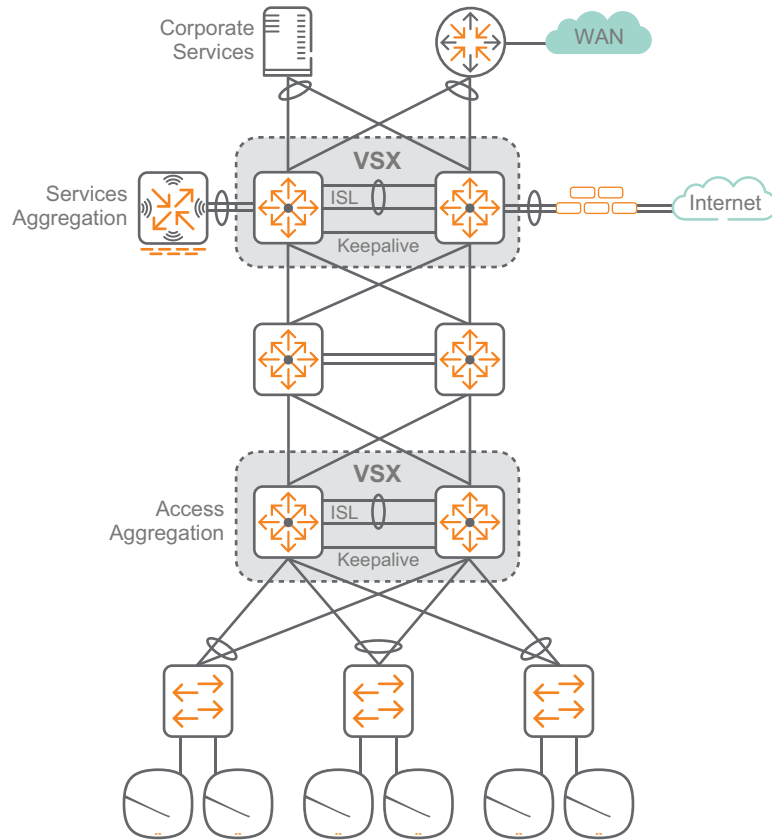
The benefits of multi-chassis link aggregation are as follows:

- **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 backplane or VSF stack consists of multiple physical devices, it is managed as a single device with a single configuration, which simplifies network design and management. Two switches using VSX appear to their neighbors as a single switch at layer-2 but as two switches at layer-3, with the key aspects of their configuration synced automatically between them.

Virtual Switching Extension

Aruba VSX is a virtualization technology for aggregation and core switches running ArubaOS-CX. It was designed to use the best features of existing HA technologies such as MC-LAG and VSF. VSX enables a distributed and redundant architecture that is highly available during upgrades, which is inherent in its architecture. The feature lets the pair of switches appear as one virtualized switch in critical areas of your network design. The configuration synchronization option allows you to sync key aspects of the primary switch to the secondary switch, which maintains operational changes across the two switches.

Figure 17 VSX in the aggregation of a three-tier design



1121A

VSX virtualizes the control plane of two switches to function as one device at layer-2 and as independent devices at layer-3. From a data-path perspective, each device does an independent forwarding lookup to decide how to handle traffic. Some of the forwarding databases, such as the MAC and ARP tables, are synchronized between the two devices via the VSX control plane over a dedicated ISL trunk. Each switch builds the layer-3 forwarding databases independently.

Benefits of VSX in the Aggregation Layer

VSX has similar benefits as VSF, but VSX offers better HA functionality during upgrades. The following benefits are grouped by functionality:


- Control plane:
 - Separate control planes to avoid shared fate issues
 - Synchronized configuration for simplicity and easy troubleshooting
 - Independently software upgradable with near zero downtime
- Layer-2 distributed MC-LAGs (aggregation switches to access switches):
 - Loop-free layer-2 multipathing (active-active forwarding)
 - Rapid failover in less than 300ms
 - Simple configuration
 - STP not required for primary failures
- Active Gateway:
 - Active-Active first hop gateway
 - No first hop redundancy protocol overhead like VRRP/HSRP
 - Simple configuration (one command)
 - DHCP relay redundancy

IP Routing

In a large 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 large 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 a typical campus network, a single area is recommended.

The access switches have a default gateway in the management VLAN for operational access, and the VLANs are terminated at the aggregation layer switches. The Internet is accessed using a static default route originating from the DMZ firewall.



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 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.

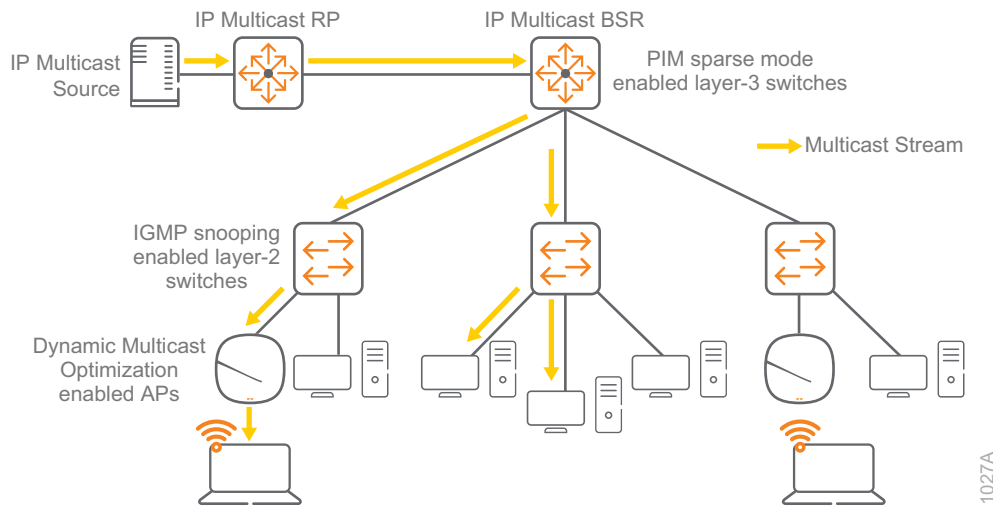
The RP is the root of the multicast tree for multicast traffic using sparse mode. Because it is the root for shared multicast traffic, the RP is normally placed at the core of the network or at the point where the most multicast senders are located. Multiple RPs can be configured for redundancy, although 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 19 IP multicast with PIM sparse mode, IGMP snooping and DMO



1027A


Core Layer

In a large LAN environment, there are often multiple aggregation layer switches. When access layer switches are located in multiple geographically dispersed buildings, you can save costly fiber-optic runs between buildings by placing an aggregation layer switch in each of those buildings. As networks grow beyond three aggregation switch pairs in a single location, organizations should add a core layer to optimize the design.

Additional aggregation layer switches are also needed when the number of access layer switches connecting to a single aggregation point exceeds the performance of the pair of aggregation switches. In a modular and scalable design, you can co-locate aggregation layer switches for data center, WAN connectivity, and Internet edge services. If your network is not large enough to warrant a standalone core, you can combine the core switch functions with the services aggregation functions using a larger modular switch for increased port capacity.

In environments where multiple aggregation layer switches exist in close proximity and where fiber optics provide high-bandwidth interconnects, a standalone core layer reduces the network complexity. The standalone core layer uses separate core switches acting independently of each other with dual equal-cost multi-path (ECMP) connections into all aggregation layer switch blocks, as shown in the following figure.

Figure 20 Standalone core layer




1054A

You can also combine the core and services functionality into a single pair of switches. When combining the core and services, we recommend VSX or stacking the two core-services switches together to allow the infrastructure devices in the services aggregation to use MC-LAG when connecting to them. This is the same

design as the access-aggregation switches and the services-aggregation switches discussed previously. If you decide you want to add a standalone core later, you can re-use the combined core-services switches as your services aggregation.

The following figure shows the core switch with a combined services aggregation in the Aruba Campus design.

Figure 21 Core and services aggregation



Core Layer Switching Features

The core layer of the LAN is a critical part of the scalable network, yet it is one of the simplest by design. The aggregation layer provides the fault and control domains and the core represents the 24x7x365 nonstop connectivity between the aggregation switch pairs.

For the fastest core layer convergence, build triangles not squares in order to take advantage of ECMP routing, which provides the best deterministic convergence. ECMP is an advanced routing strategy where next-hop packet forwarding occurs over multiple paths with identical routing metric calculations.

When considering core topologies, it is also important to use point-to-point links because link up/down changes are propagated almost immediately to the underlying protocols. Topologies with redundant ECMP links are the most deterministic and convergence is measured in milliseconds, rather than topologies that rely on indirect notification and timer-based detection, where convergence is non-deterministic and often measured in seconds.

High Performance

- The fully distributed architecture of the core switch provides up to 19.2 Tbps switching capacity with up to 7.142 billion packets per second of throughput because all switching and routing is performed in the line modules.
- A scalable system design provides investment protection in order to support future technologies and higher-speed connectivity.

Resiliency and High Availability

- The redundant and load-sharing fabrics, fan assemblies, and power supplies increases total performance and power availability while providing hitless, stateful failover.
- Hot-swappable modules allow replacement of hardware without impacting other modules.
- Separate data and control paths keeps service processing isolated, which increases security and performance of the core devices.
- ECMP enables multiple equal-cost point-to-point links in a high-speed routed environment to increase link redundancy and scale the bandwidth between devices.

Campus Wired LAN Design Summary

The Aruba Campus wired LAN provides network access for employees, APs, and IoT devices. The campus LAN also becomes the core for interconnecting the WAN, data center, and Internet access, making it a critical part of the network. The simplified access, aggregation, and core 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 reduces complexity while improving recovery times during network failures.
- The aggregation and core layers provide IP routing using OSPF and IP multicast using PIM sparse mode with redundant BSRs and RPs.
- The services aggregation is the logical place to connect critical networking devices such as corporate servers, WAN routers, and Internet-edge firewalls.
- The core is a high-speed dual-switch interconnect that provides path redundancy and sub-second failover for non-stop forwarding of packets.
- Combining the core and services aggregation into a single layer allows the network to scale when a stand-alone core is not required.

WIRED DESIGN COMPONENTS

The wired LAN in the Aruba 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. Although 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 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 for APs, IP phones, and IoT devices

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 HPE 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 SFP+ 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 and either 4 SFP+ 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 or a 24 port Smart Rate model capable of 1 and 2.5 GbE on all ports. Both PoE+ (30W) and 802.3bt (60W) or high power PoE options are available to drive current and future PoE devices. 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 8300—The Aruba 8300 series provides up to 6.4Tbps of capacity in various fixed 1U and 2U models. This switch is ideal for organizations that need to aggregate many access switches and either need or are planning for higher speed uplinks such as 10, 25 and 40 GbE at high density. This switch is also recommended for organization that have a small server farm at their location and may pair the 8300 series with other 8300s deployed as server farm top-of-rack switches. The 8300 series includes the following features:

- Intelligent monitoring and visibility with Aruba Networks Analytics Engine
- ArubaOS-CX automation and programmability using built-in APIs and Python scripts
- Advanced layer-2/3 feature set includes BGP, OSPF, VRF, active gateway, QoS, IPv6, and dynamic VXLAN with BGP-EVPN
- High availability with VSX, redundant power supplies, and fans
- Scalable line-rate interfaces at 1, 10, 25, 40, and 100 GbE for wired and wireless aggregation

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. The 5400R chassis includes the following features:

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

Aruba 3810M—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. The 3810 includes the following features:

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

Core Switches

The core layer provides high-speed routing to the aggregation blocks using a pair of redundant switches. The two switches are dual connected to all aggregation layer devices using an ECMP routing strategy where next-hop packet forwarding can occur over multiple paths that have the same routing metric. You can use ECMP with most routing protocols, including OSPF, because it is a per-hop decision.

Core Layer Switching Options

The 8400/8320 switches support the following features:

- Advanced layer-2/3 feature set includes QoS, BGP, OSPF, VRF, VRRP and IPv6
- Intelligent monitoring and visibility with Aruba Network Analytics Engine

Aruba 8400—The Aruba 8400 chassis supports line rate 10GbE/40GbE/100GbE port density, very low latency, and scalability for support of full Internet routes. The switch supports up to 256 10GbE (SFP/SFP+), or 64 40GbE (QSFP+), or 48 ports 40/100GbE (QSFP28) with eight slots for line modules and provides up to 19.2 Tbps switching capacity with up to 7.142 billion packets per second of throughput.

- Carrier-class high availability with redundant management, power and fabric

Aruba 8320—The Aruba 8320 compact switch supports line rate 10GbE/40GbE port density and very low latency. The switch supports up to 48 10GbE (SFP/SFP+), or 32 40GbE (QSFP+) and provides up to 2.5 Tbps switching capacity with up to 1905 million packets per second of throughput.

- High availability with redundant power supplies and fans

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 may need to connect servers, firewalls or other network appliances directly to the core 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

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

Deploying the Aruba Campus

The Aruba 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.

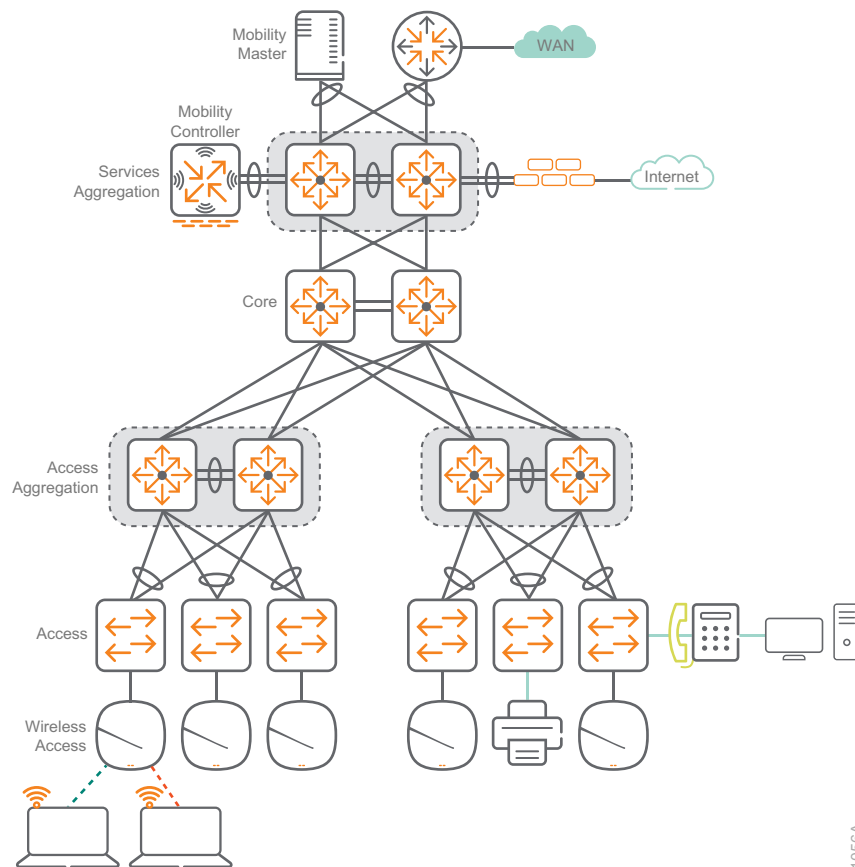
The design provides the following benefits:

- 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 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.

The following figure shows an overview of the Aruba Campus design for 500 to 10,000 users.

Figure 22 Aruba 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 500 to 10,000 users, the wired LAN has a core layer, an aggregation layer, and an access layer. With the Aruba design, the trunks between the layers use multiple active links forwarding traffic for a higher-performance network while reducing the complexity involved in traditional redundant campus 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 two VLANs, one for management and one for employee wired. The VLANs for the employee wireless and guest wireless move into the services aggregation because the wireless traffic is tunneled at layer-2 to the centralized controllers. 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 ArubaOS-Switch 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 Security Features

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

The following figure shows the wired access switch location in the Aruba Campus design.

Figure 23 Aruba 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 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 switch, 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 following figure. 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.

In a stack configuration, the port numbers incorporate both the stack member ID and the physical port number. For ease of operations, the port number order can follow the physical racking order. It is recommended that you select the desired port numbering during the initial setup, because changing them after the stack is operational will require a reload and reconfiguration of the access ports.

Figure 24 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 4 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 5 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 both have uplink ports.

Follow the following steps to connect new switches and statically assign their roles. If the switches are already configured, you should reset them to factory defaults.

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 in the order that follows the desired port numbering sequence. The first port on the first switch that you power on will become 1/1, and the first port on the second switch that you power on will become 2/1.

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



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

```
show stacking
```

```
...
```

ID	Mac Address	Model	Pri	Status
1	ecebb8-6b05c0	Aruba JL320A 2930M-24G-PoE+ Switch	230	Standby
*2	ecebb8-6bc680	Aruba JL320A 2930M-24G-PoE+ Switch	250	Commander
3	8030e0-d36a40	Aruba R0M68A 2930M-24SR-PoE-Class6...	128	Member

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



Step 4: Following the guidelines in Table 4 and Table 5, 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   ecebb8-6b05c0        Aruba JL320A 2930M-24G-PoE+ Switch 230 Standby
 *2   ecebb8-6bc680        Aruba JL320A 2930M-24G-PoE+ Switch 250 Commander
  3   8030e0-d36a40        Aruba R0M68A 2930M-24SR-PoE-Class6... 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. VSF stacking allows for longer physical separation between switches because the connectivity does not require stack cables.

The VSF stack can have as many as eight 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. Each VSF link can contain multiple physical interfaces.

The following tables show the configuration details for VSF stacking.

Table 6 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 7 VSF stacking for three-member or more 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 has only two switches, pick either switch for the Commander role because they both have uplink ports.

Follow the following steps 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. VSF physical ports must have a default configuration.



Step 1: Following the guidelines in Table 6 and Table 7, 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 VSF 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 VSF 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 16.

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

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: Configure password storage in SHA-256 on switch

```
password non-plaintext-sha256
```

Step 5: Set the idle timeout for device access to 3600 seconds (1 hour).

```
console idle-timeout 3600
```

Step 6: Enable the SSH for inbound connections.

```
ip ssh
```

Step 7: Enable the secure copy protocol (SCP).

```
ip ssh filetransfer
```

Step 8: For increased security, turn off telnet server in order to only allow inbound SSH connections.

```
no telnet-server
```

Step 9: Configure a login banner.

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

Step 10: Configure the domain name and domain name servers.

```
ip dns domain-name example.local  
ip dns server-address priority 1 8.8.8.8  
ip dns server-address priority 2 8.8.4.4
```

Step 11: 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 10.2.120.40 iburst  
ntp enable
```

Step 12: 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 13: Configure a certificate and enable 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 14: For additional security, turn off plaintext HTTP management.

```
no web-management plaintext
```

Step 15: 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] n

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 16: 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
[passwordLimited]
snmpv3 user NetAdminR auth sha [passwordRO] priv aes [passwordRO]
```

Step 17: For additional security, remove the SNMP server community public and the SNMPv3 “initial” user from the configuration.

```
no snmp-server community public
no snmpv3 user initial
```

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 for redundancy and increasing uplink capacity. By default, the uplink trunks use source and destination IP addresses to load-balance traffic between the physical interfaces. Unidirectional link detection (UDLD) monitors and disables a port when a fiber failure occurs in one direction. If a VLAN is not specified in the link-keepalive command, the UDLD packets are sent untagged.

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

```
trunk B24,D24 trk11 lacp
```

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

```
interface B24,D24 link-keepalive
link-keepalive interval 70
link-keepalive retries 6
```

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

```
spanning-tree mode rapid-pvst
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 Aruba Campus design.

Table 8 Access switch—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 A1-A24,B1-B23,C1-C8,D1-D23 (all non-trunk ports) 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 8, configure the VLAN.

Example: Management VLAN

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

Example: Wired VLAN

```
vlan 20
  name Wired
  untagged A1-A24,B1-B23,C1-C8,D1-D23
  tagged Trk11
  no ip address
  exit
```

Example: Anti-VLAN Hopping VLAN with no IP Address

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

Step 2: 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, you add the access VLAN, which you previously configured for management, to the device profile that the switches apply to traffic received from a connected AP.

The device profile in this design is used to apply the untagged VLAN command to the port where the AP is connected. The untagged VLAN is used by the AP to communicate with other APs and the mobility controllers.

On each access switch, perform the following steps.

Step 1: Configure the device profile name.

```
device-profile name "Aruba-AP-Profile"  
  untagged-vlan 10
```

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

Next, you enable multicast IGMP snooping for the layer-2 access switches.

On each access switch, perform the following step:

Step 1: Configure multicast IGMP snooping.

Example: Wired VLAN

```
vlan 20 ip igmp
```

Repeat this step for each of your VLANs where you want to send or receive multicast traffic.

1.8 Configure Access-switch Security Features

Next, you configure recommended security features 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 client VLANs and trust the trunk interface.

```
dhcp-snooping
dhcp-snooping vlan 10 20
dhcp-snooping trust trk11
```

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

```
dhcpv6-snooping
dhcpv6-snooping vlan 10 20
dhcpv6-snooping trust trk11
```

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

```
arp-protect
arp-protect vlan 20
arp-protect trust trk11
```

Step 4: Enable IP source guard globally.

```
ip source-lockdown
```

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

```
ipv6 ra-guard ports ethernet A1-A24 , B1-B23 , C1-C8 , D1-D23
```


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 shuts 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 have to trust ports with these types of devices.



```
spanning-tree A1-A24,B1-B23,C1-C8,D1-D23 bpdu-protection
spanning-tree bpdu-protection-timeout 60
```

Step 7: (Optional) Configure loop protection on the range of non-trunk ports.

```
loop-protect
loop-protect A1-A24,B1-B23,C1-C8,D1-D23
```

Wired Aggregation

The access-aggregation layer provides connectivity for the access switches and connects to the core layer using ECMP uplinks. The service-aggregation layer provides connectivity to the external networks in the campus and connects to the core layer using ECMP uplinks. The aggregation switches are layer-3 and utilize OSPF for the routing protocol.

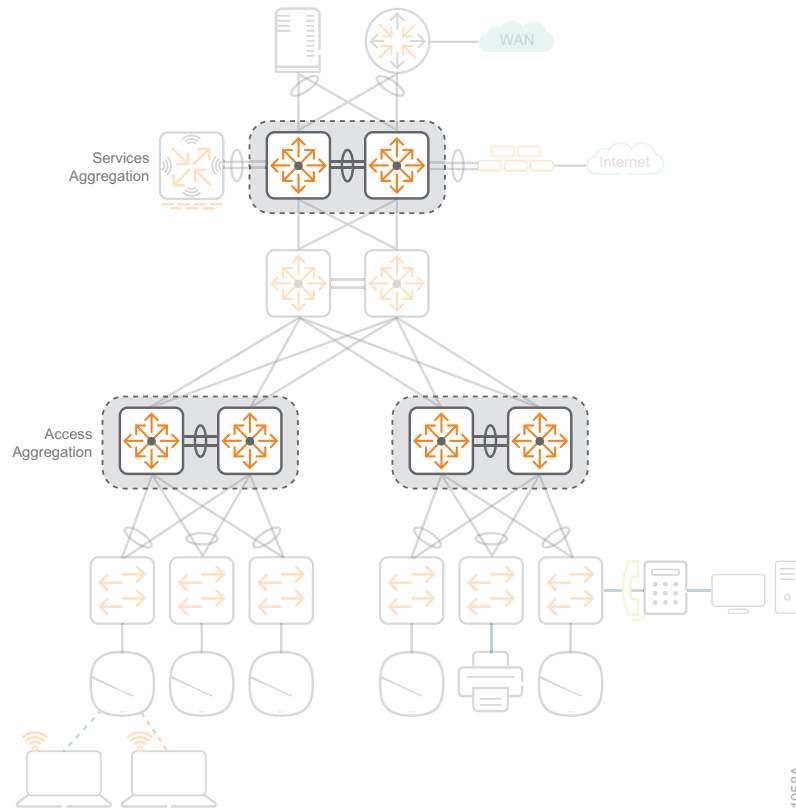
Procedures

Configuring the ArubaOS-Switch 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 Multicast IGMP Snooping
- 2.6 Configure OSPF Routing
- 2.7 Configure IP Multicast Routing

Use this section for the aggregation layer and repeat it for each aggregation switch running ArubaOS-Switch software. This includes the Aruba 5400R, Aruba 3810M, Aruba 2930M, and Aruba 2930F switches. You can use this section for standalone switches, switch stacks, or modular aggregation switches. If you have a switch running ArubaOS-CX in your aggregation layer, skip to the next section of the guide. The following figure shows the access and services aggregation switch with ArubaOS-Switch locations in the Aruba Campus design.

Figure 25 Aruba Campus design—access and services aggregation with ArubaOS-Switch



1058A

2.1 Configure Aggregation-switch Stacking

Stacking allows multiple 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 following table shows the configuration details for backplane stacking.

Table 9 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 following steps 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 in the order that follows the desired port numbering sequence. The first port on the first switch that you power on will become 1/1, and the first port on the second switch that you power on will become 2/1.

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



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

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



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. VSF stacking allows for longer physical separation between switches because the connectivity does not require stack cables.

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. Each VSF link can contain multiple physical interfaces.

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 following table shows the configuration details for VSF stacking.

Table 10 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 following steps 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

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: Configure password storage in SHA-256 on switch.

```
password non-plaintext-sha256
```

Step 5: Set the idle timeout for device access to 3600 seconds (1 hour).

```
console idle-timeout 3600
```

Step 6: Enable SSH for inbound connections.

```
ip ssh
```

Step 7: Enable SCP.

```
ip ssh filetransfer
```

Step 8: For increased security, turn off telnet server in order to only allow inbound SSH connections.

```
no telnet-server
```

Step 9: Configure a login banner.

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

Step 10: Configure the domain name and domain name servers.

```
ip dns domain-name example.local  
ip dns server-address priority 1 8.8.8.8  
ip dns server-address priority 2 8.8.4.4
```

Step 11: Configure the 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 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 10.2.120.40 iburst
ntp enable
```

Step 12: 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 13: 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 14: For additional security, turn off plaintext HTTP management.

```
no web-management plaintext
```


Step 15: 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] n

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 16: 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
[passwordLimited]
snmpv3 user NetAdminR auth sha [passwordRO] priv aes [passwordRO]
```

Step 17: For additional security, remove the SNMP server community public and initial user from the configuration.

```
no snmp-server community public
no snmpv3 user initial
```

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 uses 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 trk21 lacp
```

Repeat this step for each trunk.

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

```
int 1/A1,2/A1 link-keepalive  
link-keepalive interval 70  
link-keepalive retries 6
```

Repeat this step for each set of uplink ports.

Step 3: 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 mode rapid-pvst  
spanning-tree priority 0
```

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

```
logging severity info
```

2.4 Configure Aggregation-switch VLANs

Next, you configure the VLANs for the aggregation switch. The aggregation switch is the default gateway for the user or services VLANs. The non-trunk ports are configured as untagged in the wired or data center VLANs. The uplink trunk ports from the access layer are configured as tagged for the user VLANs and untagged for VLAN 777.

The uplink ports into the core layer are configured as point-to-point for ECMP routing. They use a 30-bit mask because each subnet only needs two IP addresses. The aggregation switch has four separate connections into the core to allow two redundant paths into each of the standalone core switches. The uplink ports can

be individual physical interfaces or lag interfaces that use the link aggregation control protocol (LACP) to combine two or more physical ports into a single trunk interface. In the examples shown below, single ports are used between the aggregation and core layers. 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 accepts the first offer it receives.

The following tables provide the VLAN assignments for the Aruba Campus design. If you are configuring an access aggregation switch, use the information from the first table. If you are configuring a services aggregation switch, use the information from the second table.

Table 11 Access-aggregation switch—VLAN assignments, IP addresses, and tagging

VLAN name	VLAN ID	IP address	IP helper address	Tagged/untagged ports
Management	10	10.2.20.1/22	10.2.120.40	Tagged Trk21-Trk25
Wired	20	10.2.24.1/22	10.2.120.40	Untagged 1/1-1/48 (all non-trunk ports) Tagged Trk11-Trk15
Core 1-1	216	10.2.202.18/30	N/A	Untagged 1/C7 (core switch 1, port 1)
Core 1-2	220	10.2.202.22/30	N/A	Untagged 2/C7 (core switch 1, port 2)
Core 2-1	224	10.2.202.26/30	N/A	Untagged 1/C8 (core switch 2, port 1)
Core 2-2	228	10.2.202.30/30	N/A	Untagged 2/C8 (core switch 2, port 2)
Anti-VLAN hopping	777	N/A	N/A	Untagged Trk21-Trk25

Table 12 Services-aggregation switch—VLAN assignments, IP addresses, and tagging

VLAN name	VLAN ID	IP address	IP helper address	Tagged/untagged ports
Wireless	330	10.2.8.1/22	10.2.120.40	Tagged Trk11-Trk12 (mobility controllers)
Guest	340	10.2.12.1/22	10.2.120.40	Tagged Trk11-Trk12(mobility controllers)
Data Center	120	10.2.120.1/24	N/A	Untagged 1/1-1/48 (all non-trunk ports)
Internet DMZ	140	10.2.140.1/24	N/A	Tagged 1/47-1/48 (firewall)
Mobility Controllers	160	10.2.160.1/24	N/A	Tagged 1/A1-1/A2,2/3-2/8 (mobility controllers)
Core 1	600	10.2.206.1/30	N/A	Untagged Trk1 (core switch 1, lag 1)
Core 2	604	10.2.206.5/30	N/A	Untagged Trk2 (core switch 2, lag 2)

On each aggregation switch, perform the following steps:

Step 1: Configure the aggregation VLANs.

Example: Management VLAN with Helper Address

```
vlan 10
  name "Management"
  tagged Trk21-Trk25
  ip address 10.2.0.1 255.255.252.0
  ip helper-address 10.2.120.40
  exit
```

Example: Core VLAN without Helper Address

```
vlan 216
  name "Core 1-1"
  untagged 1/C7
  ip address 10.2.202.18 255.255.255.252
  exit
```

Example: Anti-VLAN Hopping VLAN without IP Address and Helper Address

```
vlan 777
  name Anti-VLAN hopping
  untagged Trk21-Trk25
  exit
```

Repeat this step for each VLAN that matches the type of switch you are configuring.

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.

Example: Management, Wired, and Anti-VLAN Hopping VLANs

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

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

Example: Management VLAN

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

2.5 Configure Multicast IGMP Snooping

Next, you enable multicast IGMP snooping for the service-aggregation switch because IGMP snooping must be enabled for DMO to work on the APs. You do not need to configure IGMP snooping on the access-aggregation switches.

On the service aggregation switch, perform the following step:

Step 1: Configure multicast IGMP snooping.

Example: Wireless VLAN

```
vlan 20 ip igmp
```

Repeat this step for each of your VLANs where you want to send or receive multicast traffic.

2.6 Configure OSPF Routing

In this procedure, you configure OSPF as the layer-3 routing protocol. This design uses area backbone (0.0.0.0) for the entire campus network. Configure the user and services VLANs as passive, because there are no devices that need routing protocol updates attached to them.

On each aggregation switch, perform the following steps:

Step 1: Configure the loopback interface.

```
interface loopback 1
  ip address 10.2.255.50
```

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 10.2.255.50
```

Step 3: Configure OSPF. Use nonstop routing with modular aggregation switches.

```
router ospf
  area backbone
  nonstop
  enable
  exit
```

Step 4: Configure the loopback interface.

```
interface loopback 1
  ip ospf 10.2.255.50 area backbone
```

Step 5: Configure the uplinks to core. Use one VLAN per uplink.

```
vlan 216
  untagged 1/C7
  ip address 10.2.202.18 255.255.255.252
  ip ospf 10.2.202.18 area backbone
  ip ospf network-type point-to-point
```

Repeat this step for each core VLAN.

Step 6: Configure the user or services VLANs with passive mode.

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

Repeat this step for each user or services VLAN with an IP address.

2.7 Configure IP Multicast Routing

Next, you enable multicast routing for the layer-3 aggregation switches. The design is based on sparse mode multicast operation.

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

```
ip multicast-routing
```

Step 2: Enable PIM.

```
router pim
  enable
```

Step 3: Configure PIM sparse mode on the VLANs and allow any IP address to source multicast streams.

Example: Wired VLAN on access-aggregation

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

Example: Wireless VLAN on service-aggregation

```
vlan 30
  ip pim-sparse
  ip-addr any
  exit
exit
```

Repeat this step for each of your VLANs where you want to send or receive multicast traffic.

Procedures

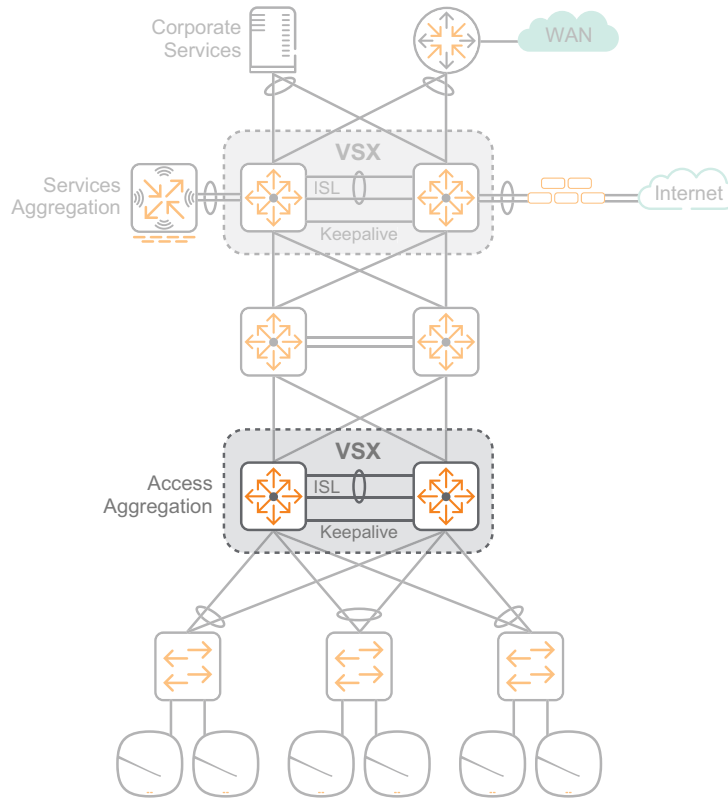
Configuring the ArubaOS-CX Aggregation Switch

- 3.1 Configure the VSX on the Aggregation-switch
- 3.2 Configure the Aggregation-switch Base Features
- 3.3 Configure Uplink Ports from Aggregation to Access
- 3.4 Configure Aggregation-switch VLANs
- 3.5 Configure OSPF Routing
- 3.6 Configure IP Multicast

Use this section for the aggregation layer and repeat it for each wired aggregation switch running ArubaOS-CX software. This includes the Aruba 8300 and 8400 series switches. You can use this section for a standalone switch or a pair of VSX configured switches. If you do not have a switch running ArubaOS-CX in your aggregation layer, skip to the next section, "Campus Wireless LAN."

The following figure shows the wired aggregation with ArubaOS-CX location in the Aruba Campus design.

Figure 26 Aruba Campus design—access aggregation with ArubaOS-CX



1122A

3.1 Configure the VSX on the Aggregation-switch

VSX virtualizes the control plane of two switches, which allows them to function as one device at layer-2 and as independent devices at layer-3. From a data-path perspective, each device performs its own forwarding lookup to decide how to handle traffic. We recommend two switches with VSX for the aggregation layer.

This design uses a LAG interface for the VSX Inter-Switch Link (ISL) connection between the switches with at least two physical interfaces. IP addresses are not needed on this interface because the ISL protocol is layer-2. Set the MTU to the maximum size allowed on the interfaces.

The following table shows the VSX ISL VLAN and LAG assignments for this design.

Table 13 VSX ISL VLAN and LAG assignments

VLAN description	VLAN ID	LAG	Trunk native	Trunk allowed	MTU
VSX ISL LAG	1	128	1	all	9198

This design uses a single physical interface for the keepalive direct-connection between the switches. The interface is placed into a VRF to isolate the routing from the global VRF routing table, which prevents other traffic from using the directly connected link. You can use the same IP subnet and addresses on all your VSX switch pairs because they are isolated by the VRF.

The following table shows the VSX keepalive VRF and IP address assignments for this design.

Table 14 VSX keepalive VRF and IP address assignments

VRF	IP address primary	IP address secondary
VSX-Keepalive	10.99.99.1/30	10.99.99.2/30

On each aggregation switch, perform the following steps.

Step 1: Configure the ISL LAG interface between the two switches. Sync the VLANs between the switches. Select the native VLAN and allow all VLANs to be trunked. Enable LACP mode active.

```
interface lag 128
  no shutdown
  no routing
  vlan trunk native 1 tag
  vlan trunk allowed all
  lacp mode active
```

Step 2: Configure at least two ISL physical interfaces between the two switches. Set the MTU to 9198.

```
interface 1/1/53
  no shutdown
  mtu 9198
  lag 128
interface 1/1/54
  no shutdown
  mtu 9198
  lag 128
```

Step 3: Configure a keepalive VRF to create an isolated network between the two switches.

```
vrf VSX-Keepalive
```

Step 4: Configure the keepalive physical interface between the two switches. Attach the keepalive VRF to the interface. Configure an IP subnet that is not used anywhere else in your network, so it is easily identified.

```
interface 1/1/51  
  no shutdown  
  vrf attach VSX-Keepalive  
  description VSX Keepalive  
  ip address 10.99.99.1/30
```

Step 5: Configure VSX. Define a common system-mac for L2 protocols. Make one switch primary and the other switch secondary. Use the keepalive interface IP addresses and VRF as the peer and source address. Select the configuration items you want VSX to sync between the two switches.

```
vsx  
  system-mac 00:00:10:02:55:30  
  inter-switch-link lag 128  
  role primary  
  keepalive peer 10.99.99.2 source 10.99.99.1 vrf VSX-Keepalive  
  vsx-sync dns lldp mclag-interfaces ssh stp-global time vsx-global
```

Step 6: For the other switch in the VSX pair, repeat this procedure using the appropriate values.

3.2 Configure the Aggregation-switch Base Features

On each aggregation switch, perform the following steps.

Step 1: Configure the switch host name.

```
hostname Aggregation-Switch
```

Step 2: Configure the unrestricted administrator password.

```
user admin password plaintext [password]
```

Step 3: Require a username and password for console access using local credentials.

```
aaa authentication login console local
```

Step 4: Set the idle timeout for device access to 60 minutes (1 hour).

```
cli-session  
  timeout 60
```

Step 5: Enable SSH server for inbound connections in the default vrf.

```
ssh server vrf default
```

Step 6: Configure a login banner.

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

Step 7: Configure the domain name and domain name servers.

```
ip dns domain-name example.local  
ip dns server-address 8.8.8.8  
ip dns server-address 8.8.4.4
```

Step 8: Configure the network time protocol (NTP) with time zone and daylight savings time.

```
clock timezone pst8pdt  
ntp enable  
ntp server 10.2.120.40 iburst
```

Step 9: If the date on your device is not current, use the **clock date** command to set the date to today's date.

```
clock date YYYY-MM-DD
```

Step 10: Configure HTTP Secure (HTTPS) server for web access.

```
https-server vrf default
```

Step 11: Configure SNMP server in the default vrf.

```
snmp-server vrf default
```

Step 12: Configure SNMP server community to override the default name **public**.

```
snmp-server community NetAdminPriv
```

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

```
snmpv3 user NetAdminRW auth sha auth-pass plaintext [passwordRW] priv aes priv-  
pass plaintext [passwordRW]  
snmpv3 user NetAdminLimited auth sha auth-pass plaintext [passwordLimited] priv  
aes priv-pass plaintext [passwordLimited]  
snmpv3 user NetAdminRO auth sha auth-pass plaintext [passwordRO] priv aes priv-  
pass plaintext [passwordRO]
```

3.3 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 use source and destination IP addresses to load-balance traffic between the physical interfaces.

On each aggregation switch, perform the following steps.

Step 1: Configure the multi-chassis lag interface with lacp mode active and enable the interface.

```
interface lag 11 multi-chassis
    no shutdown
    no routing
    lacp mode active
```

Configure the physical interfaces for the dynamic lag group and enable UDLD. Configure the UDLD retries to 6.

```
interface 1/1/1
    no shutdown
    lag 11
    udld
    udld retries 6
```

Repeat this step for each uplink interface in the lag group on both switches.

3.4 Configure Aggregation-switch VLANs

The layer-3 aggregation switch is the default gateway for the user VLANs, and they need an IP address. The uplink lag interfaces are configured with VLAN 777 as native and the user VLANs as allowed.

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.

You can mitigate double tagging by creating an unused VLAN that is configured only 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 accepts the first offer it receives.

The following table provides the VLAN assignments for the Aruba Campus design.

Table 15 VLAN assignments, IP addresses, and tagging

VLAN description	VLAN ID	IP address for agg 1	IP address for agg 2	IP helper address	Active gateway IP	Active gateway MAC address
Management	10	10.2.0.2/22	10.2.0.3/22	10.2.120.40	10.2.0.1	00:00:10:02:00:01
Wired	20	10.2.4.2/22	10.2.4.3/22	10.2.120.40	10.2.4.1	00:00:10:02:04:01
Anti-VLAN hopping	777	N/A	N/A	N/A	N/A	N/A

On each aggregation switch, perform the following steps.

Step 1: Configure the aggregation VLAN and interface.

Example: Management VLAN for Aggregation 1

```
vlan 10
interface vlan 10
description Management
ip address 10.2.0.2/22
ip helper-address 10.2.120.40
```

Example: Anti-VLAN Hopping VLAN

```
vlan 777
name Anti-VLAN hopping
```

Repeat this step for each VLAN in the previous table.

Step 2: (Optional) If you are using VSX for your aggregation switches, configure the active gateway IP and MAC addresses on all your user VLANs. The virtual MAC address must be unique, so matching it to the IP address is an easy way to keep it simple.

Example: Management VLAN

```
interface vlan 10
active-gateway ip 10.2.0.1 mac 00:00:10:02:00:01
```

Repeat this step for each user VLAN in the previous table.

Step 3: Enable Rapid-PVST for all VLANs. 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 10,20,777
spanning-tree vlan 10,20,777 priority 0
spanning-tree
```

3.5 Configure OSPF Routing

In this procedure, you configure OSPF as the layer-3 routing protocol. This design uses area backbone (0.0.0.0) for the entire network. Use the router loopback IP address as the OSPF router ID. Configure the user and services VLANs as passive, because there are no devices that need routing protocol updates attached to them.

Perform this procedure on each core switch.

Step 1: Configure the loopback interface.

```
interface loopback 1
ip address 10.2.255.30/32
```

Step 2: Configure OSPF.

```
router ospf 1
router-id 10.2.255.30
passive-interface default
area 0.0.0.0
enable
```

Step 3: Configure the interface for OSPF.

Example: Loopback Interface

```
interface loopback 1
  ip ospf 1 area 0.0.0.0
```

Example: Physical Interface

```
interface 1/1/1
  ip ospf 1 area 0.0.0.0
  ip ospf network point-to-point
```

Example: Lag Interface

```
interface lag 1
  ip ospf 1 area 0.0.0.0
  ip ospf network point-to-point
```

Example: Wired VLAN Interface (Passive on Non-user Interfaces)

```
interface vlan20
  ip ospf 1 area 0.0.0.0
  ip ospf passive
```

Repeat this step for each active interface.

3.6 Configure IP Multicast

In this procedure, you enable multicast routing for the aggregation switch. The design is based on sparse-mode multicast operation. You use BSRs and RPs to provide a simple yet scalable way to provide a highly resilient RP environment. Make the aggregation switches the primary and secondary BSR and RP candidates, because they are in the middle of the network and all multicast traffic must pass through them anyway.

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.

Do not use the interfaces between the switches as the source IP interfaces because if one of the switches goes down, the adjacent port on the other switch also goes down. We recommend you use the loopback interface as the source for both the BSR and RP.

If there are multiple PIM-SM devices on a LAN, a DR must be elected to avoid duplicating multicast traffic for connected hosts. The PIM device with the highest IP address becomes the DR for the LAN unless you force the DR election using the **ip pim dr-priority** command.

Perform the following steps on each aggregation switch.

Step 1: Configure PIM sparse mode on the interfaces with IP addresses where you want to send multicast traffic. Set the DR priority value on the interface of the primary switch of the VSX pair. The highest priority on a given LAN segment will be elected as the DR.

Example: Physical Interface of Primary VSX Switch

```
interface 1/1/49
  ip pim-sparse enable
  ip pim-sparse dr-priority 10
```

Example: Employee VLAN Interface of Primary VSX Switch

```
interface vlan20
  ip pim-sparse enable
  ip pim-sparse dr-priority 10
```

Repeat this step for each interface and VLAN with an IP address where you want to send multicast traffic.

Step 2: Enable PIM and configure the switch as a BSR candidate by using a source IP interface of the loopback interface, and then select a priority that makes one of them higher than the other.

```
router pim
  enable
  bsr-candidate source-ip-interface loopback1
  bsr-candidate priority 60
```

Step 3: Configure the switch as a candidate RP by using a source IP interface pointing at the access switch, a group prefix of 224.0.0.0/4, and then select a priority that makes one of them lower than the other.

```
rp-candidate source-ip-interface loopback1
rp-candidate group-prefix 224.0.0.0/4
rp-candidate priority 40
```

Step 4: Save the configuration to flash.

```
write memory
```


Wired Core

The core layer provides high-speed layer-3 connectivity for the aggregation layer switches. It can also provide services aggregation functions when needed. The decision to use a standalone core layer depends on the number of aggregation layer switches and if your services are combined in a single location or spread across several aggregation blocks. With this Aruba Campus architecture, you can start with a combined core and services design, and then migrate to a standalone core when needed. The ECMP uplinks between the access aggregation switches, and the core layer remains the same with either model.

Use this section for the core layer and repeat it for each core switch. This section describes configuring both standalone core switches and core switches with a combined services aggregation. Choose the group of procedures that match the type of core network you are deploying.

Procedures

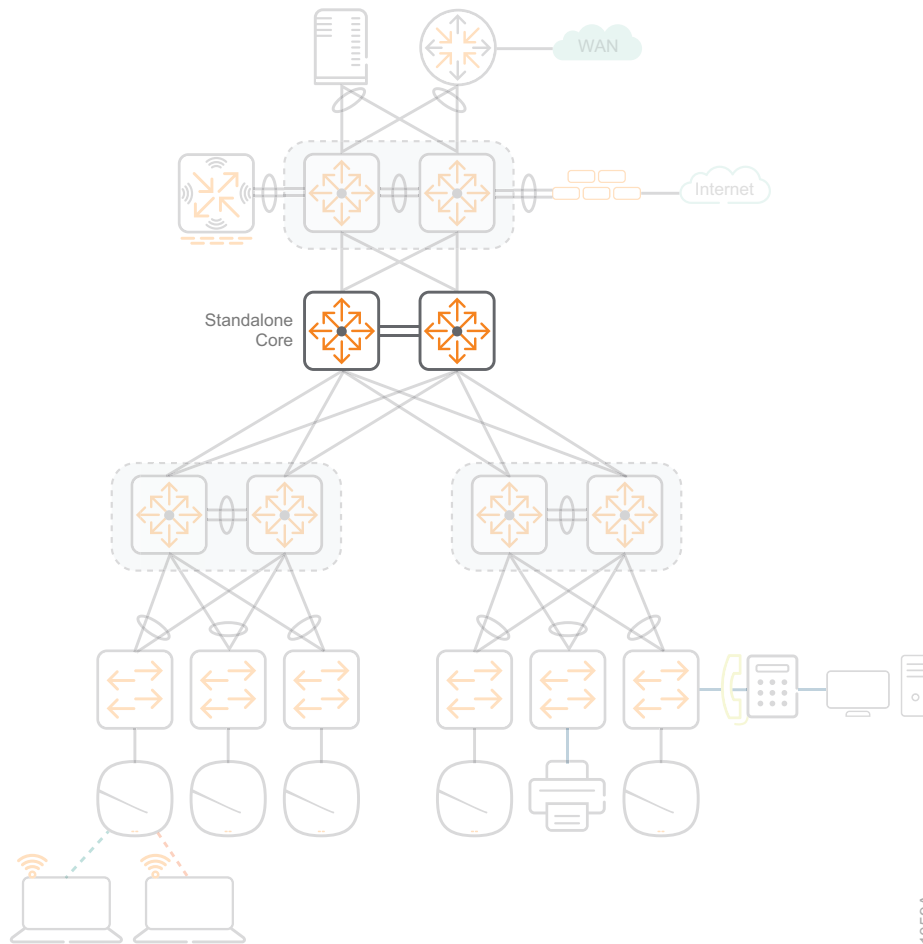
Configuring the Standalone Core Switches

- 4.1 Configure the Core-switch Base Features
- 4.2 Configure Uplink Ports from Core to Aggregation
- 4.3 Configure OSPF Routing
- 4.4 Configure IP Multicast

The standalone core switches do not use stacking technologies to combine them into a single logical switch. Each standalone core switch operates on its own and uses ECMP routing techniques to load-balance traffic to and from the aggregation layer switches.

The following figure shows the standalone core switches in the Aruba Campus design.

Figure 27 Aruba Campus design—standalone core



1059A

4.1 Configure the Core-switch Base Features

The switch has two levels of access: administrator and operator. The administrator 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.

Complete this procedure on each core switch.

Step 1: Configure the switch host name.

```
hostname Core-Switch
```

Step 2: Configure the unrestricted administrator password.

```
user admin password
Changing password for user admin
Enter password: [passwordAdmin]
Confirm new password: [passwordAdmin]
```

Step 3: Enable SSH server for inbound connections in the default vrf.

```
ssh server vrf default
```

Step 4: Configure a login banner.

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

Step 5: Configure the domain name and domain name servers.

```
ip dns domain-name example.local
ip dns server-address 8.8.8.8
ip dns server-address 8.8.4.4
```

Step 6: Configure the network time protocol (NTP) with time zone and daylight savings time.

```
clock timezone pst8pdt
ntp server 10.2.120.40 iburst
```

Step 7: If the date on your device is not current, use the **clock date** command to set the date to today's date.

```
clock date YYYY-MM-DD
```

Step 8: Configure HTTP Secure (HTTPS) server for web access.

```
https-server vrf default
```

Step 9: Configure SNMP server in the default vrf.

```
snmp-server vrf default
```

Step 10: Configure SNMP server community to override the default name **public**.

```
snmp-server community NetAdminPriv
```

Step 11: Create full read-write, limited read-write, and read-only users for SNMPv3.

```
snmpv3 user NetAdminRW auth sha auth-pass plaintext [passwordRW] priv aes priv-  
pass plaintext [passwordRW]  
snmpv3 user NetAdminLimited auth sha auth-pass plaintext [passwordLimited] priv  
aes priv-pass plaintext [passwordLimited]  
snmpv3 user NetAdminRO auth sha auth-pass plaintext [passwordRO] priv aes priv-  
pass plaintext [passwordRO]
```

4.2 Configure Uplink Ports from Core to Aggregation

The procedure configures the uplink ports from the core switch to the aggregation switch. The uplink ports are point-to-point for ECMP routing. They use a 30-bit mask because each subnet only needs two IP addresses. The uplink ports can be individual physical interfaces or lag interfaces that use LACP to combine two or more physical ports into a single trunk interface. By default, the LACP trunks use source and destination IP addresses to load-balance traffic between the physical interfaces.

This procedure describes how to configure both individual physical and lag interfaces. Repeat the appropriate options below for each uplink interface.

Option 1: Using Individual Physical Interfaces

Step 1: Configure the physical interface.

```
interface 1/2/1  
  no shutdown  
  ip address 10.2.202.1/30
```

Step 2: Repeat this procedure for each uplink interface.

Option 2: Using Lag Interfaces

Step 1: Configure the lag interface.

```
interface lag 1  
  no shutdown  
  ip address 10.2.206.2/30  
  lacp mode active
```

Step 2: Configure the physical interfaces for the dynamic lag group.

```
interface 1/1/31
  no shutdown
  lag 1
interface 1/1/32
  no shutdown
  lag 1
```

Step 3: Repeat this procedure for each uplink interface.

4.3 Configure OSPF Routing

Next, you configure OSPF as the layer-3 routing protocol. This design uses area backbone (0.0.0.0) for the entire campus network. Use the router loopback IP address as the OSPF router ID. Redistribute the connected and static routes.

Perform this procedure on each core switch.

Step 1: Configure the loopback interface.

```
interface loopback 1
  ip address 10.2.255.50/32
```

Step 2: Configure OSPF.

```
router ospf 1
  router-id 10.2.255.50
  area 0.0.0.0
  enable
```

Step 3: Configure the interface for OSPF.

Example: Loopback interface

```
interface loopback 1
  ip ospf 1 area 0.0.0.0
```

Example: Physical interface

```
interface 1/2/1
  ip ospf 1 area 0.0.0.0
  ip ospf network point-to-point
```

Example: Lag interface

```
interface lag 1
  ip ospf 1 area 0.0.0.0
  ip ospf network point-to-point
```

Repeat this step for each interface.

4.4 Configure IP Multicast

In this procedure, you enable multicast routing for the core switch. The design is based on sparse-mode multicast operation. You use BSRs and RPs to provide a simple yet scalable way to provide a highly resilient RP environment. Make the core switches the primary and secondary BSR and RP candidates, because they are in the middle of the network and all multicast traffic must pass through them anyway.

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.

Do not use the interfaces between the switches as the source IP interfaces because if one of the switches goes down, the adjacent port on the other switch also goes down. We recommend you use the interfaces that point to the services aggregation layer as the source for both the BSR and RP.

Perform the following steps on each core switch.

Step 1: Configure PIM sparse mode on the interfaces with IP addresses.

Example: Physical Interface

```
interface 1/2/1
    ip pim-sparse enable
```

Example: Lag Interface

```
interface lag 1
    ip pim-sparse enable
```

Repeat this step for each active interface.

Step 2: Enable PIM and configure the switch as a BSR candidate by using a source IP interface pointing at the services aggregation switch, and then select a priority that makes one of them higher than the other.

```
router pim
    enable
    bsr-candidate source-ip-interface 1/2/1
    bsr-candidate priority 60
```

Step 3: Configure the switch as a candidate RP by using a source IP interface pointing at the services aggregation switch, a group prefix of 224.0.0.0/4, and then select a priority that makes one of them lower than the other.

```
rp-candidate source-ip-interface 1/2/1
rp-candidate group-prefix 224.0.0.0/4
rp-candidate priority 40
```

Configuring the Core Switches with a Combined Services Aggregation


If you are planning to combine the core and services functionality into a single pair of switches, we recommend stacking them together to allow the infrastructure devices in the services aggregation to use MC-LAG when connecting to the switches. This is the same design as the access-aggregation switches and the services-aggregation switches discussed previously. If you decide you want to add a standalone core later, you can re-use the combined core-services switches as your services aggregation. For configuration guidance, refer to the aggregation layer procedures.

Note If the number of ARP entries in your environment exceeds 4000 wireless clients, we recommend using a standalone core and a separate services aggregation block as outlined in the previous procedure.



The following figure shows the core switch with a combined services aggregation in the Aruba Campus design.

Figure 28 Aruba Campus design—core and services aggregation



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 a Mobility Master pair, mobility controller cluster, and campus APs with employee and guest WLANs and VLANs. The APs use a management VLAN to communicate between each other and the mobility controllers. The employee traffic is tunneled from the employee SSID to the employee VLAN and the guest traffic from the guest SSID to the guest VLAN on the uplink ports of the mobility controllers. A policy created in the mobility controller allows the employees to access the entire network, while the guests can access the guest VLAN, DHCP server, DNS service, and HTTP/HTTPS, in order to access web sites on the Internet.


Procedures

Configuring the Mobility Master

- 5.1 Configure the Mobility Master System Setup
- 5.2 Configure Mobility Master Redundancy
- 5.3 Configure Mobility Master Database Synchronization
- 5.4 Install and Enable Licenses

Use this section to configure the Mobility Master. The following figure shows the Mobility Master pair in the services aggregation of the Aruba Campus design. They also are commonly deployed in the virtual server environment of the data center.

Figure 29 Aruba Campus design—Mobility Master pair



1061A

The Aruba Mobility Master acts as a single point of configuration for global policies such as firewall rules, authentication, and RF to simplify the administration and maintenance of a wireless network. The design uses a centralized, multi-tier architecture that provides a clear separation between management, control, and forwarding functions. The Mobility Master provides the management and control, while the mobility controllers provide the forwarding.

The Mobility Master runs ArubaOS 8 and centrally manages the cluster of mobility controllers. The Mobility Master consolidates all-master, single master-multiple local, and multiple master-local deployments from previous versions of ArubaOS into a single centralized deployment model.

You can deploy the Mobility Master as a virtual machine or on a hardened appliance. For more information about the virtual appliance installation, see the [ArubaOS 8.5.0.0 Virtual Appliance Installation Guide](#).


The information from the following table includes the Mobility Master virtual router ID, IP addresses, priority and VLANs used in the procedures below.

Table 16 Example Mobility Master virtual router ID, VLAN ID, IP addresses and priority

Name	Virtual router ID	Virtual IP address	Priority	VLAN ID	VLAN name	Local IP address	Peer IP address
Example-MM1	120	10.2.120.100	200	120	AMS-Office-SC-MM	10.2.120.80	10.2.120.90
Example-MM2	120	10.2.120.100	100 (default)	120	AMS-Office-SC-MM	10.2.120.90	10.2.120.80

For physical wired redundancy, it is recommended you configure the Mobility Master virtual machines on two different host machines. Each virtual machine has a single connection into the virtual switch and the virtual switch is dual-connected to two physical interfaces on the host machine. The host machines are dual-connected to two different physical switches in your services aggregation with NIC teaming using LACP, as depicted in the following figure.

Figure 30 Mobility Master physical wired redundancy



5.1 Configure the Mobility Master System Setup

Step 1: After the wired network is fully operational, connect a pair of Mobility Masters to two different physical switches in the services aggregation switch stack and power them on. The initial power-on sequence takes several minutes to complete.

Step 2: Open a console session into the **primary Mobility Master** and enter the following values in the setup dialog.

```
Enter System name [ArubaMM]: Example-MM1
Enter Controller VLAN ID [1]: 120
Enter Controller VLAN port [GE 0/0/0]:
Enter Controller VLAN port mode (access|trunk) [access]:
Enter VLAN interface IP address [172.16.0.254]: 10.2.120.80
Enter VLAN interface subnet mask [255.255.255.0]:
Enter IP default gateway [none]: 10.2.120.1
Do you wish to configure IPV6 address on vlan (yes|no) [yes]: no
Enter Country code (ISO-3166), <ctrl-I> for supported list: us
You have chosen Country code US for United States (yes|no)?: yes
Enter the controller's IANA Time zone [America/Los_Angeles]:
Enter Time in UTC [14:59:52]: 22:04:00
Enter Date (MM/DD/YYYY) [6/10/2019]:
Enter Password for admin login (up to 32 chars): [password]
Re-type Password for admin login: [password]
```

Current choices are:

```
System name: Example-MM1
Controller VLAN id: 120
Controller VLAN port: GE 0/0/0
Controller VLAN port mode: access
VLAN interface IP address: 10.2.120.80
VLAN interface subnet mask: 255.255.255.0
IP Default gateway: 10.2.120.1
Option to configure VLAN interface IPV6 address: no
Country code: us
IANA Time Zone: America/Los_Angeles
```

If you accept the changes the switch will restart!

Type <ctrl-P> to go back and change answer for any question

Do you wish to accept the changes (yes|no) **yes**

Step 3: For the standby Mobility Master, repeat the previous step changing the variables as required.

5.2 Configure Mobility Master Redundancy

Next, you configure redundancy for the Mobility Master pair. To maintain a highly redundant network, the administrator configures a second Mobility Master to act as a hot standby for the primary Mobility Master using the Virtual Router Redundancy Protocol (VRRP).

When the master is unavailable, the standby becomes the master and takes ownership of the virtual IP address. All network elements (APs and other controllers) are configured to access the virtual IP address, thereby providing a transparent redundant solution.

The priority level of the VRRP instance is used in the election mechanism for the master. The highest priority value becomes the VRRP master. The default priority value is 100, so use a higher number on the primary Mobility Master.

Note For Mobility Master virtual appliances in a VMware vSphere environment, promiscuous mode needs to be turned on for VRRP to work. For security reasons, we recommend a separate virtual machine group in vSwitch for ports using VRRP and that you only allow promiscuous mode in that group.



Step 1: Browse to the **primary Mobility Master**, enter the following information, and then click **Log in**.

- Username—**admin**
- Password—**[password]**

Step 2: Click the icon on the top left of the page to expand the menu and navigate to **Mobility Master > Example-MM1 > Configuration > Redundancy > L2 redundancy > Virtual Router Table**, and then click **+**.

The screenshot shows the configuration interface for L2 redundancy. At the top, there are two tabs: "L2 redundancy" (selected) and "L3 redundancy". Below the tabs, there is a section titled "Virtual Router Table" with a dropdown arrow. Underneath, there is a table with the following columns: "ROUTER NAME", "IPV4 ADDRESS", "IPV6 ADDRESS", "VLAN", "ADMIN STATE", and a grid icon. The table is currently empty. At the bottom left of the table area, there is a blue plus sign (+) icon.

Step 3: In the **New Virtual Router** window at the bottom of the page, enter the following information, and then click **Submit**.

- ID—**120**
- IP version—**IPv4**
- Authentication password—**[password]**
- Retype authentication password—**[password]**
- IP address—**10.2.120.100** (virtual IP address)
- Priority—**200** (this is the primary Mobility Master)
- Enable router pre-emption—**Enabled**
- Admin state—**UP**
- VLAN—**120** (VLAN from the Mobility Master system setup)
- Leave the rest of the fields blank

Step 4: At the top right of the page, click **Pending Changes > Deploy changes > Close**.

Step 5: Navigate to **Mobility Master > Example-MM1 > Configuration > Redundancy > L2 redundancy > Master Redundancy**, enter the following information, and then click **Submit**.

- Master VRRP—**120** (virtual router ID from previous step)
- IP address of peer—**10.2.120.90** (IP address of standby Mobility Master)
- Authentication—**IPSec Key** (this field and the following two will not show up until you enter an IP address of the peer)
- IPSec key of peer—**[password]**
- Retype IPSec key—**[password]**

Step 6: Click **Pending Changes > Deploy changes > Close**.

Step 7: Navigate to **Mobility Master > Example-MM1 > Configuration > Interfaces > VLANs**, and then click **+**.



Step 8: In the **New VLAN** window, enter the following information, and then click **Submit**.

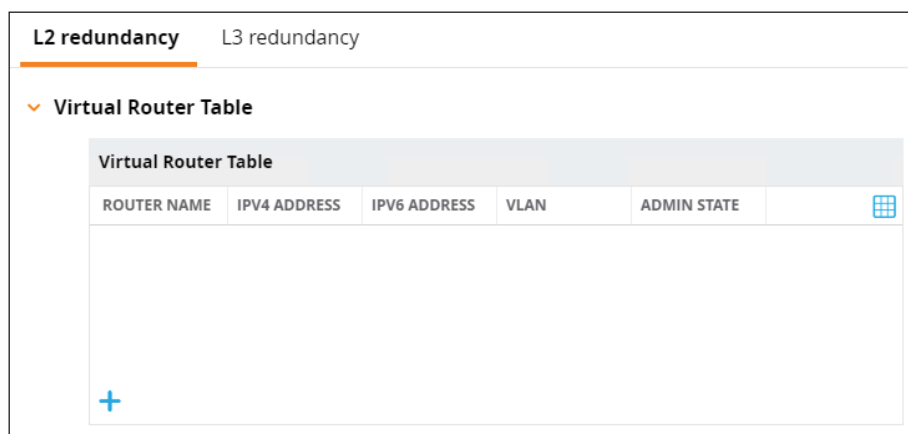
- VLAN name—**AMS-Office-SC-MM**
- VLAN ID/Range—**120**

Step 9: Click **Pending Changes > Deploy changes > Close**.

Step 10: Browse to the **standby Mobility Master**, enter the following information, and then click **Log in**.

- Username—**admin**
- Password—**[password]**

Step 11: Click the icon on the top left of the page to expand the menu and navigate to **Mobility Master > Example-MM2 > Configuration > Redundancy > L2 redundancy > Virtual Router Table**, and then click **+**.



Step 12: In the **New Virtual Router window**, enter the following information, and then click **Submit**.

- ID—**120**
- IP version—**IPv4**
- Authentication password—**[password]**
- Retype authentication password—**[password]**
- IP address—**10.2.120.100** (virtual IP address)
- Priority—**100** (this is the standby Mobility Master)
- Admin state—**UP**
- VLAN—**120**
- Leave the rest of the fields blank

Step 13: Click **Pending Changes > Deploy changes > Close**.

Step 14: Navigate to **Mobility Master > Example-MM2 > Configuration > Redundancy > Master Redundancy**, enter the following information, and then click **Submit**.

- Master VRRP—**120** (virtual router ID from previous step)
- IP address of peer—**10.2.120.80** (IP address of primary Mobility Master)
- Authentication—**IPSec Key** (this field and the following two will not show up until you enter an IP address of the peer)
- IPSec key of peer—**[password]**
- Retype IPSec key—**[password]**

Step 15: Click **Pending Changes > Deploy changes > Close**.

Step 16: Navigate to **Mobility Master > Example-MM2 > Configuration > Interfaces > VLANs**, and then click **+**.



Step 17: In the **New VLAN** window, enter the following information, and then click **Submit**.

- VLAN name—**AMS-Office-SC-MM**
- VLAN ID/Range—**120**

Step 18: Click **Pending Changes > Deploy changes > Close**.

5.3 Configure Mobility Master Database Synchronization

Next, you configure the synchronization between the redundant Mobility Master databases. The standby Mobility Master must synchronize its database from the primary Mobility Master after the APs are communicating with the mobility controllers over a secure channel. This ensures that all certificates, IPsec keys, and campus AP whitelist entries are available to the backup master.

You should also synchronize the database immediately when APs are added or removed from the campus AP whitelist to ensure that the backup master has the latest settings.

Note The database synchronization must be configured on the top-level Mobility Master folder and not at the individual device level.



Step 1: Login to the primary Mobility Master and navigate to **Mobility Master > Configuration > Redundancy > Master Redundancy**, enter the following information, and then click **Submit**.

- Database synchronization slider—**Enabled**
- Sync period—**60** (default)

L2 redundancy

> Virtual Router Table

▼ **Master Redundancy**

Database synchronization:

Sync period: Minutes

Step 2: Click **Pending Changes > Deploy changes > Close**.

5.4 Install and Enable Licenses

Next, you install and enable ArubaOS licenses. To learn more about licenses and licensing features, see the [ArubaOS 8.5.0.x Licensing Guide](#).

Step 1: Navigate to **Mobility Master > Configuration > License > License Inventory**, and then on the top right of the page, click **+**.

License management: Aruba Support Portal (ASP) External license server Manual

License Usage		License Inventory			+
LICENSE	DESCRIPTION	STATUS	EXPIRATION	INSTALLED to this Mobility Master	
AP	Access Points	⊖ Not Licensed	Not Licensed	0	
PEFNG	Policy Enforcement Firew...	⊖ Not Licensed	Not Licensed	0	
RFP	RF Protect(WIP,Spectrum,...	⊖ Not Licensed	Not Licensed	0	
ACR	Advanced Cryptography	⊖ Not Licensed	Not Licensed	0	
WebCC	Web Content Classification	⊖ Not Licensed	Not Licensed	0	
MM	Mobility Master Virtual A...	⊖ Not Licensed	Not Licensed	0	
MC-VA-RW	Controller Virtual Applia...	⊖ Not Licensed	Not Licensed	0	
MC-VA-EG	Controller Virtual Applia...	⊖ Not Licensed	Not Licensed	0	
MC-VA-II	Controller Virtual Annlia...	⊖ Not Licensed	Not Licensed	0	

Step 2: Paste your previously obtained licenses into the **Install Licenses** window, and then click **OK** and **Submit**.

Step 3: Click **Pending Changes > Deploy changes > Close**.

Step 4: Navigate to **Mobility Master > Configuration > License > License Usage**, and then click **Global License Pool**.

Step 5: In the **Usage for Global License Pool** window, enable the following features, and then click **Submit**.

- PEF—**Enable**
- RF Protect—**Enable**

Usage for Global License Pool			
	AP	PEF	RF Protect
Feature Enabled	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Scope	Per-AP	Per-AP	Per-AP
Pool Size	50	50	50
Expired Licenses	0	0	0
Actual Pool Size	50	50	50
Licenses Used	0	0	0
Licenses Remaining Available	50	50	50

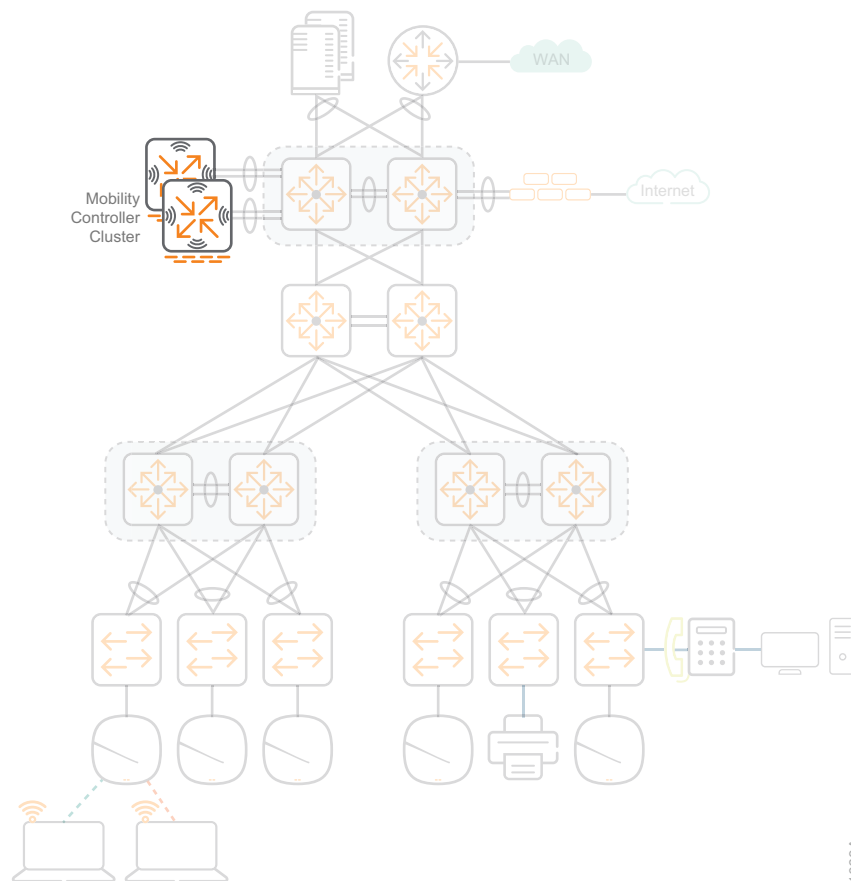
Step 6: Click **Pending Changes > Deploy changes > Close**.

Configuring the Mobility Controller Cluster

- 6.1 Initial Mobility Controller Setup
- 6.2 Configure the Hierarchy
- 6.3 Adopt Mobility Controllers in the Mobility Master
- 6.4 Add Adopted Mobility Controllers to Local Group
- 6.5 Configure LACP Interfaces
- 6.6 Configure Mobility Controller Clustering
- 6.7 Configure Mobility Controller Cluster VRRP for AP Provisioning

Use this section to configure the mobility controller cluster. The following figure shows the mobility controller cluster in the services aggregation of the Aruba Campus design.

Figure 31 Aruba Campus design—mobility controller cluster



1063A

The mobility controllers can be deployed as virtual machines or as dedicated hardware devices. To learn more about the virtual appliance installation, please refer to [ArubaOS 8.5.0.0_Virtual Appliance Installation Guide](#).

6.1 Initial Mobility Controller Setup

Next, you configure the initial system setup for the mobility controller. The information from the following table includes the IP addresses and VLANs used in the procedures below.

Table 17 Example mobility controllers IP addresses and VLAN ID

Name	IP address	Default gateway	VLAN ID	VLAN name	Master switch IP address (VRRP)
CLS1-MC-1	10.2.160.10/24	10.2.160.1	160	AMS-Office-SC-MC	10.2.120.100
CLS1-MC-2	10.2.160.20/24	10.2.160.1	160	AMS-Office-SC-MC	10.2.120.100

Step 1: After the Mobility Masters are operational, power on your mobility controllers.

The initial power-on sequence takes several minutes to complete.

Step 2: From a console session, enter the following values in the setup dialog.

```
Auto-provisioning is in progress. It requires DHCP and Activate servers
Choose one of the following options to override or debug auto-provisioning...
  'enable-debug'      : Enable auto-provisioning debug logs
  'disable-debug'    : Disable auto-provisioning debug logs
  'mini-setup'       : Start mini setup dialog. Provides minimal
customization and requires DHCP server
  'full-setup'       : Start full setup dialog. Provides full customization
  'static-activate'  : Provides customization for static or PPPOE ip
assignment. Uses activate for master information
```

```
Enter Option (partial string is acceptable): full-setup
```

```
Are you sure that you want to stop auto-provisioning and start full setup
dialog? (yes/no): yes
```

```
Enter System name [Aruba7205_02_F5_40]: CLS1-MC-1
```

```
Enter Switch Role (master|standalone|md) [md]:
```

```
Enter IP type to terminate IPsec tunnel or secured websocket connection
(ipv4|ipv6) [ipv4]:
```

Enter Master switch IP address/FQDN or ACP IP address/FQDN: **10.2.120.100**
Enter Master switch Type? (MM|ACP) [MM]:
Is this a VPN concentrator for managed device to reach Master switch (yes|no) [no]:
This device connects to Master switch via VPN concentrator (yes|no) [no]:
Is Master switch Virtual Mobility Master? (yes|no) [yes]:
Master switch Authentication method (PSKwithIP|PSKwithMAC) [PSKwithIP]:
Enter IPsec Pre-shared Key: **[password]**
Re-enter IPsec Pre-shared Key: **[password]**
Do you want to enable L3 Redundancy (yes|no) [no]:
Enter Uplink Vlan ID [1]: **160**
Enter Uplink port [GE 0/0/0]: **GE 0/0/4**
Enter Uplink port mode (access|trunk) [access]: **trunk**
Enter Native VLAN ID [1]:
Enter Uplink Vlan IP assignment method (dhcp|static) [static]:
Enter Uplink Vlan Static IP address [172.16.0.254]: **10.2.160.10**
Enter Uplink Vlan Static IP netmask [255.255.255.0]:
Enter IP default gateway [none]: **10.2.160.1**
Enter DNS IP address [none]: **10.2.120.50**
Do you wish to configure IPV6 address on vlan (yes|no) [yes]: **no**
Do you want to configure dynamic port-channel (yes|no) [no]:
This controller is restricted, please enter country code (US|PR|GU|VI|MP|AS|FM|MH) [US]:
You have chosen Country code US for United States (yes|no)?: **yes**
Enter the controller's IANA Time zone [America/Los_Angeles]:
Enter Time in UTC [14:59:52]: **22:04:00**
Enter Date (MM/DD/YYYY) [6/10/2019]:
Do you want to create admin account (yes|no) [yes]:
Enter Password for admin login (up to 32 chars): **[password]**
Re-type Password for admin login: **[password]**

Current choices are:

```
System name: CLS1-MC-1
Switch Role: md
IP type to terminate IPsec tunnel or secured websocket connection: ipv4
Master switch IP address or FQDN: 10.2.120.100
Is this VPN concentrator: no
Connect via VPN concentrator: no
IPsec authentication method: PSKwithIP
Vlan id for uplink interface: 160
Uplink port: GE 0/0/4
Uplink port mode: trunk
Native VLAN id: 1
Uplink Vlan IP assignment method: static
Uplink Vlan static IP Address: 10.2.160.10
Uplink Vlan static IP net-mask: 255.255.255.0
Uplink Vlan IP default gateway: 10.2.160.1
Domain Name Server to resolve FQDN: 10.2.120.50
Option to configure VLAN interface IPV6 address: no
Country code: US
IANA Time Zone: America/Los_Angeles
Admin account created: yes
```

Note: These settings require IP-Based-PSK configuration on Master switch

```
If you accept the changes the switch will restart!
Type <ctrl-P> to go back and change answer for any question
Do you wish to accept the changes (yes|no)yes
```

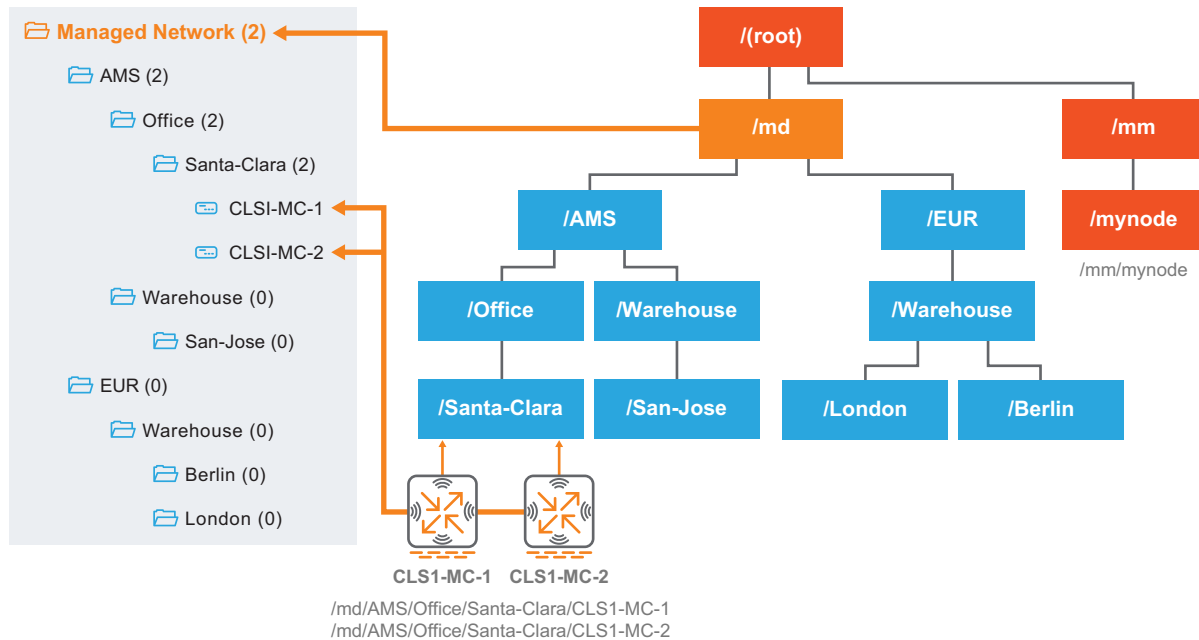
Step 3: For each additional mobility controller, repeat Step 2, changing the variables as required.

6.2 Configure the Hierarchy

Next, you configure the hierarchical groups. For more information about Mobility Master configuration hierarchy, see Chapter 1 of the [ArubaOS 8.5.0.x User Guide](#).

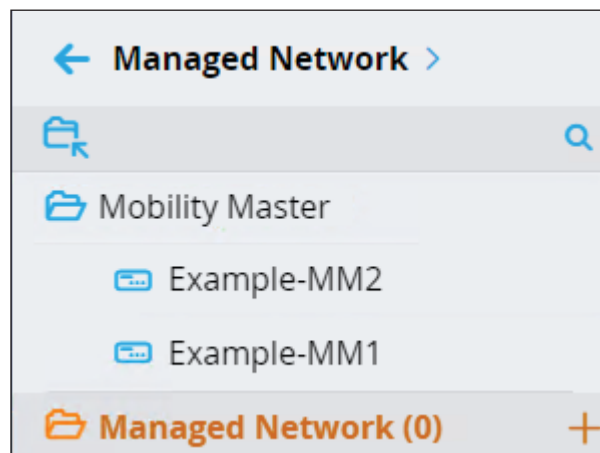
The following is an example of the configuration hierarchy used in this guide with the corresponding menu hierarchy shown on the right-hand side.

Figure 32 Example configuration hierarchy



1123A

Step 1: Login to the **primary Mobility Master**, navigate to **Managed Network**, and then on the right side, click **+**.

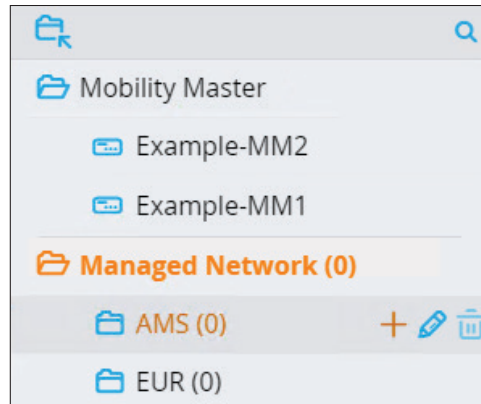


Step 2: In the **Add** window, enter the following information, and then click **Submit**.

- Select—**Group**
- Name—**AMS**

Step 3: For each additional level-1 group, repeat Step 1 and Step 2, changing the variable as required.

Step 4: After completing you level-1 groups, navigate to **Managed Network > AMS**, and then click +.




Step 5: In the **Add** window, enter the following information, and then click **Submit**.

- Select—**Group**
- Name—**Office**

Step 6: For each additional level-2 group, repeat Step 4 and Step 5, changing the variable as required.

Step 7: After completing you level-2 groups, navigate to **Managed Network > AMS > Office**, and then click +.



Step 8: In the **Add** window, enter the following information, and then click **Submit**.


- Select—**Group**
- Name—**Santa-Clara**

Step 9: For each additional level-3 group, repeat Step 7 and Step 8, changing the variable as required.

6.3 Adopt Mobility Controllers in the Mobility Master

Next, you configure the Mobility Master to adopt mobility controllers into the system.

Step 1: Navigate to **Mobility Master > Configuration > Controllers > Local Controller IPsec Keys**, and then on the lower left side click **+**.



IPV4 ADDRESS OF THE LOCAL	IPV6 ADDRESS OF THE LOCAL	KEY	MAC ADDRESS OF THE LOCAL	CERT TYPE	

Step 2: In the **Add New IPsec Controller** window, enter the following information, and then click **Submit**.

- Authentication—**IPsec Key** (default)
- Local controller IPV4—**10.2.160.10** (IP address of mobility controller)
- IPsec key—**[password]**
- Retype IPsec key—**[password]**

Step 3: For each additional mobility controller, repeat Step 1 and Step 2, changing the variables as required.

Step 4: Click **Pending Changes > Deploy changes > Close**.

6.4 Add Adopted Mobility Controllers to Local Group

Step 1: To obtain the MAC address of the mobility controllers, SSH into the **primary Mobility Master** and issue the **show switches** command.

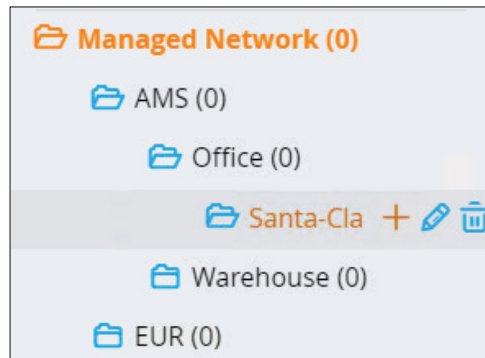
Note It will take several minutes for the mobility controllers to establish their secure connections to the Mobility Master and for them to appear in the **show switches** command.



The MAC address and associated IP address for each mobility controller can be found under the **Configuration State** column as shown in the partial output of the command below.

```
(Example-MM1) [mynode] #show switches
All Switches
IP Address  IPv6 Address  Name          Location      Type      Model      Version      Status
ID          ID
-----
10.2.120.80  None          Example-MM1   Building1.Floor1  master   ArubaMM-VA  8.5.0.3_72498  up
10.2.120.90  None          Example-MM2   Building1.Floor1  standby  ArubaMM-VA  8.5.0.3_72498  up
10.2.160.10  None          CLS1-MC-1    Building1.Floor1  MD       Aruba7205   8.5.0.3_72498  up
10.2.160.20  None          CLS1-MC-2    Building1.Floor1  MD       Aruba7205   8.5.0.3_72498  up
Total Switches:4
(Example-MM1) [mynode] #
```

Step 2: Navigate to **Managed Network > AMS > Office > Santa-Clara** and then click **+**.



Step 3: In the **Add** window, enter the following information, and then click **Submit**.

- Select—**Controller**
- Hostname—**CLS1-MC-1**
- MAC address— **20:4c:03:02:f5:40** (from the show switches command)
- Type—**A7205** (must match the type of controller)

Step 4: For each additional mobility controller, repeat Step 2 and Step 3, changing the variables as required.

Step 5: Click **Pending Changes > Deploy changes > Close**.


6.5 Configure LACP Interfaces

Optional

This optional procedure configures LACP interfaces on the mobility controllers. If you have virtual mobility controllers or you do not plan to use LACP interfaces in your environment, skip to the next procedure.

For physical wired redundancy, it is recommended you configure the mobility controllers with LACP interfaces and connect them to two different physical switches in your services aggregation, as depicted in the following figure.

Figure 33 Mobility controller physical wired redundancy



To make it easier to identify the ports where your mobility controllers are connected to your LAN switches, it is recommended you turn on LLDP in the mobility controllers before enabling LACP.

Step 1: Navigate to **Managed Network > AMS > Office > Santa-Clara > CLS-MC-1 > Configuration > Interfaces > Ports**, and then click **GE-0/0/4**.

Ports							
PORT	ADMIN STA...	TRUSTED	POLICY	MODE	NATIVE VL...	ACCESS VL...	TRUNK VLA...
GE-0/0/0	Enabled	--	Not-defined	access	1	1	1-4094
GE-0/0/1	Enabled	--	Not-defined	access	1	1	1-4094
GE-0/0/2	Enabled	--	Not-defined	access	1	1	1-4094
GE-0/0/3	Enabled	--	Not-defined	access	1	1	1-4094
GE-0/0/4	Enabled	✓	Not-defined	trunk	1	1	1-4094

Step 2: At the bottom of the **GE-0/0/4** window, click **Show advanced options**, scroll down to the bottom again, enter the following information, and then click **Submit**.

- LLDP Transmission—**Enable** (slider bar to the right)
- Transmit Interval—**30** seconds (default)
- Transmit hold—**4** seconds (default)
- Fast Transmit Interval—**1** second (default)
- Fast Transmit hold—**4** seconds (default)
- LLDP Reception—**Enable**

LLDP Transmission:	<input checked="" type="checkbox"/>
Transmit Interval:	<input type="text" value="30"/> sec
Transmit hold:	<input type="text" value="4"/> sec
Fast Transmit Interval:	<input type="text" value="1"/> sec
Fast Transmit hold:	<input type="text" value="4"/> sec
LLDP Reception:	<input checked="" type="checkbox"/>
LLPD-MED:	<input type="checkbox"/>

Step 3: For each additional LLDP interface on each mobility controller, repeat Step 1 and Step 2, choosing the correct interface as required.

Step 4: Click **Pending Changes > Deploy changes > Close**.

Step 5: Navigate to **Managed Network > AMS > Office > Santa-Clara > CLS1-MC-1 > Configuration > Interfaces > Ports**, and then in the **Port Channel** section, click **+**.

Step 6: From the **New Port Channel** window, enter the following information, and then click **Submit**.

- ID—**PC-0** (choose the first available channel)

New Port Channel	
ID:	<input type="text" value="PC-0"/> ▼
<input type="button" value="Cancel"/> <input type="button" value="Submit"/>	

Step 7: From **PC-0** section, enter the following information.

- Protocol—**LACP**
- LACP mode—**active**
- Admin state—**Enable**
- Trust—**Enable**
- Mode—**Trunk**
- Native VLAN—**160**
- Allowed VLANs—click + and select from the drop down **160** (user VLANs will be added in a subsequent step)

The screenshot displays the configuration page for a Port Channel (PC-0). The configuration is as follows:

- Port channel ID: PC-0
- Protocol: LACP
- LACP mode: active
- Port members: GE-0/0/4, GE-0/0/5
- Admin state:
- Trust:
- Policy: Not-defined
- Mode: Trunk
- Native VLAN: 160
- Allowed VLANs: Allow specified VLANs

VLAN	TRUSTED / UNTRUSTED
160	Trusted

Step 8: On Port members, click **Edit**.

Step 9: Add the port numbers to be included in the port channel, and then click **Submit**.

Available/Selected Ports

<p>Available:</p> <div style="border: 1px solid gray; padding: 5px; min-height: 100px;"> <p>GE-0/0/0</p> <p>GE-0/0/1</p> <p>GE-0/0/2</p> <p>GE-0/0/3</p> </div>	<p>></p> <p>>></p> <p><</p> <p><<</p>	<p>Selected:</p> <div style="border: 1px solid gray; padding: 5px; min-height: 100px;"> <p>GE-0/0/4</p> <p style="background-color: #f4a460;">GE-0/0/5</p> </div>
<p>Cancel OK</p>		

Step 10: For each additional LACP interface on each mobility controller, repeat Step 5 through Step 7, choosing the correct interfaces as required.

Caution After you deploy the changes, the controllers will become unavailable on the network until you add matching LACP configurations on the corresponding ports on the services aggregation switch.



Step 11: Click **Pending Changes > Deploy changes > Close**.

Step 12: To obtain the corresponding LAN switch port numbers for your mobility controllers, SSH into the **services aggregation switch** and issue the **show lldp info remote-device** command.

The switch local ports for each mobility controller can be found under the **LocalPort** column as shown in the output of the command below.

```
3810-s1# sh lldp info remote-device
LLDP Remote Devices Information
-----
LocalPort | ChassisId | PortId | PortDescr | SysName
-----
1/1 | 000b86-bbba7 | GE0/0/24 | MPLS Agg | VPNC-DC1-1
1/2 | 000b86-bbfa7 | GE0/0/24 | MPLS Agg | VPNC-DC1-2
1/3 | 204c03-02f540 | GE0/0/4 | GE0/0/4 | CLS1-MC-1
1/4 | 204c03-02f6a0 | GE0/0/4 | GE0/0/4 | CLS1-MC-2
1/11 | 94f128-eeef00 | 1/1/32 | 1/1/32 | 8400-C1-1
1/12 | 94f128-ee4f00 | 1/1/32 | 1/1/32 | 8400-C1-2
1/16 | ecebb8-1774e2 | 27 | 1/A2 | Pod2-Tor-MPLS
2/1 | 000b86-bbba7 | GE0/0/25 | INET Agg | VPNC-DC1-1
2/2 | 000b86-bbfa7 | GE0/0/25 | INET Agg | VPNC-DC1-2
2/3 | 204c03-02f540 | GE0/0/5 | GE0/0/5 | CLS1-MC-1
2/4 | 204c03-02f6a0 | GE0/0/5 | GE0/0/5 | CLS1-MC-2
2/11 | 94f128-eeef00 | 1/1/31 | 1/1/31 | 8400-C1-1
2/12 | 94f128-ee4f00 | 1/1/31 | 1/1/31 | 8400-C1-2
OOBM | d06726-726cca | 112 | 3/6 | Pod1-Tor-1
OOBM | d06726-726cca | 113 | 3/7 | Pod1-Tor-1
```

Step 13: Configure the dual-port trunks with LACP.

```
trunk 1/3,2/3 trk11 lacp
```

Repeat this step for each mobility controller trunk.

Step 14: Configure the service aggregation user VLANs with the LACP trunk interfaces.

The following table provides the VLAN assignments for the user VLANs.

Table 18 Services-aggregation switch—VLAN assignments, IP addresses, and tagging

VLAN name	VLAN ID	IP address	IP helper address	Tagged/Untagged ports
Wireless	330	10.2.8.1/22	10.2.120.40	Tagged Trk11-Trk12 (mobility controllers)
Guest	340	10.2.12.1/22	10.2.120.40	Tagged Trk11-Trk12 (mobility controllers)
Mobility Controllers	160	10.2.160.1/24	N/A	Untagged Trk11-Trk12 (mobility controllers)

Example: Wireless Employee VLAN tagged

```
vlan 330
  tagged Trk11-Trk12
exit
```

Example: Mobility controller VLAN untagged

```
vlan 160
  untagged Trk11-Trk12
exit
```

Repeat this step for each user VLAN.

6.6 Configure Mobility Controller Clustering

In this procedure, you configure mobility controller layer-2 clustering. A cluster combines multiple controllers together to provide high availability and load balancing for all clients and ensures service continuity when a failover occurs.


For more information about controller clustering, see Chapter 18 of the [ArubaOS 8.5.0.x User Guide](#).

The individual mobility controller VRRP IP addresses you configure in this procedure allow authorization servers, like Aruba ClearPass or RADIUS, to make a change of authorization request for users on specific mobility controllers to the virtual IP address created below. When the mobility controller is taken out of service, the standby mobility controller servicing the user handles the requests from the authorization server. The table and figure below show the VRRP IP addresses the authentication servers use to communicate with the individual mobility controllers in the cluster.

Table 19 Mobility controller VRRP IP addresses and VLANs

Mobility controller	IP address	Multicast VLAN	VRRP IP address	VRRP VLAN
CLS-MC-1	10.2.160.10	160	10.2.160.110	160
CLS-MC-2	10.2.160.20	160	10.2.160.120	160

Figure 34 Individual mobility controller VRRP IP addresses for authentication servers



Step 1: Navigate to **Managed Network > AMS > Office > Santa-Clara > Configuration > Services > Clusters**, and then click +.

Step 2: In **New Cluster Profile** window, click +.

New Cluster Profile

Name:

Controllers

IP ADDRESS	GROUP	VRRP-IP	VRRP-VLAN	RAP PUBLIC IP	MCAST-VLAN	
<div style="display: flex; justify-content: space-between; align-items: center;"> + </div>						

Step 3: In the **Add Controller** window, enter the following information, and then click **OK**.

- IP version—**IPv4**
- IP address—**10.2.160.10** (IP address of mobility controller)
- Group—**None** (default)
- Priority—**Blank** (default)
- MCastVLAN—**160** (VLAN ID for multicast traffic)
- VRRP-IP—**10.2.160.110** (VRRP IP address of this mobility controller for authorization servers)
- VRRP-VLAN—**160** (VLAN ID for this VRRP instance)

Note For a mobility controller virtual appliance in VMware vSphere, promiscuous mode and forged retransmits need to be allowed for VRRP to work. For security reason, create a separate virtual machine group in vSwitch for ports using VRRP and only allow promiscuous mode in that group.



Step 4: For each additional mobility controller, repeat Step 2 and Step 3, changing the variables as required.

Step 5: After all the mobility controllers have been added, in the **New Cluster Profile** window, enter the following information, and then click **Submit**.

- Name—**AMS-Office-SC-Cluster**

IP ADDRESS	GROUP	VRRP-IP	VRRP-VLAN	RAP PUBLIC IP	MCAST-VLAN
10.2.160.10	--	10.2.160.110	160	--	160
10.2.160.20	--	10.2.160.120	160	--	160

Step 6: Click **Pending Changes > Deploy changes > Close**.

Step 7: Navigate to **Managed Network > AMS > Office > Santa-Clara > CLS1-MC-1 > Configuration > Services > Clusters > Cluster Profile**, enter the following information, and then click **Submit**.

- Cluster group-membership—**AMS-Office-SC-Cluster**
- Exclude VLANs—**1**

Step 8: For each additional mobility controller, repeat Step 7.

Step 9: Click **Pending Changes > Deploy changes > Close**.

6.7 Configure Mobility Controller Cluster VRRP for AP Provisioning


Next, you configure a cluster-wide VRRP for AP provisioning.

The cluster-wide VRRP IP address you configure in this procedure is for AP controller discovery. The virtual IP address is configured at the cluster level and is used in DHCP and DNS servers for dynamic discovery or it is statically defined in the Mobility Master. The table and figure below show the VRRP IP address the APs use to find the mobility controller cluster.

Table 20 Cluster-wide VRRP ID, IP address and VLAN

Cluster	VRRP ID	VRRP IP address	VRRP VLAN
AMS-Office-SC-Cluster	160	10.2.160.100	160

Figure 35 Cluster-wide VRRP IP address for AP provisioning



Step 1: Navigate to **Managed Network > AMS > Office > Santa-Clara > Configuration > Interfaces > VLANs**, and then click **+**.



Step 2: In the **New VLAN** window, enter the following information, and then click **Submit**.

- VLAN name—**AMS-Office-SC-MC**
- VLAN ID/Range—**160**

Note This step is naming the uplink VLAN already configured in the initial mobility controller setup procedure, which makes it available at the cluster level for the step below.



Step 3: Click **Pending Changes > Deploy changes > Close**.

Step 4: Navigate to **Managed Network > AMS > Office > Santa-Clara > Configuration > Redundancy > L2 redundancy > Virtual Router Table**, and then click the **+**.

Step 5: In the **New Virtual Router** window, enter the following information, and then click **Submit**.

- ID—**160**
- IP version—**IPv4**
- Authentication password—**[password]**
- Retype authentication password—**[password]**
- IP address—**10.2.160.100** (cluster-wide VRRP IP address for controller discovery)
- Admin state—**UP**
- VLAN—**160** (VLAN ID for this VRRP instance)

Note For a mobility controller virtual appliance in VMware vSphere, promiscuous mode and forged retransmits need to be allowed for VRRP to work. For security reason, create a separate virtual machine group in vSwitch for ports using VRRP and only allow promiscuous mode in that group.



Step 6: Click **Pending Changes > Deploy changes > Close**.

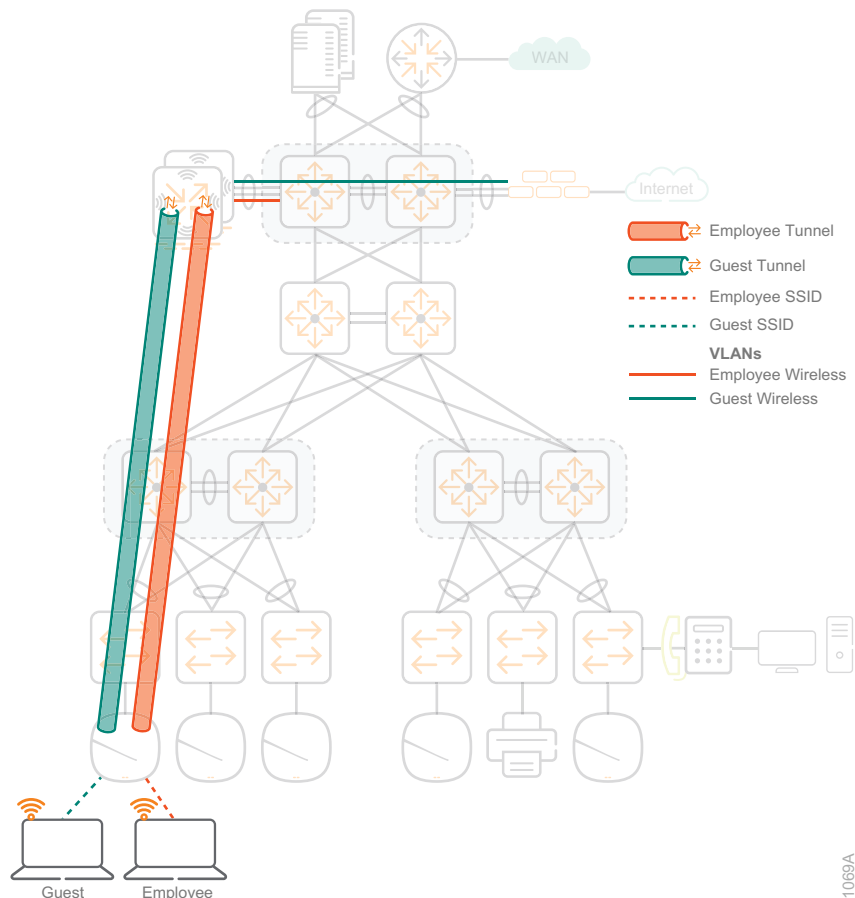
Procedures

Configuring the Wireless LANs

- 7.1 Configure External RADIUS Server
- 7.2 Configure WLAN for Employee SSID
- 7.3 Configure the Employee SSID Security
- 7.4 Configure the Employee SSID Access Rules
- 7.5 Configure WLAN for Guest SSID
- 7.6 Configure the Guest SSID Security
- 7.7 Configure Local VLANs for Employee and Guest SSIDs
- 7.8 Configure Mobility Controller Interface with VLANs
- 7.9 Configure Mobility Controller Guest VLAN with an IP Address

Use this section to configure the wireless LANs. The following figure shows wireless VLANs in the Aruba Campus design.

Figure 36 Aruba Campus design—wireless VLANs



1069A

7.1 Configure External RADIUS Server

Next, you configure the Mobility Master with an external RADIUS server that authenticates users on the WLAN. The ArubaOS software allows you to use an external authentication server or the internal user database of the Mobility Master to authenticate clients to the wireless network. For this example, the RADIUS server is configured at the Managed Network level and will be available for all mobility controllers.

Step 1: Navigate to **Managed Network > Configuration > Authentication > Auth Servers > All Servers**, and then click +.

Step 2: In the **New Server** window, enter the following information, and then click **Submit**.

- Name—**Example-CPPM**
- IP address—**10.2.120.10** (IP address of external RADIUS server)
- Type—**RADIUS**

Step 3: Navigate to **Managed Network > Configuration > Authentication > Auth Server > All Servers**, and then select **Example-CPPM**.

NAME	TYPE	IP ADDRESS / HOST...	SERVER GROUP
Example-CPPM	RADIUS	10.2.120.10	--
Internal	--	--	default internal

Step 4: In the **Server Options** window, enter the following information, and then click **Submit**.

- IP address—**10.2.120.10** (External RADIUS server added previously)
- Auth port—**1812** (default)
- Accounting port—**1813** (default)
- Shared key—**[password]**
- Retype key—**[password]**

Step 5: Click **Pending Changes > Deploy changes > Close**.

7.2 Configure WLAN for Employee SSID

You can configure a WLAN to provide different network access or services to users on the same physical network. For example, you can configure a WLAN to provide access to guest users and another WLAN to provide access to employee users through the same APs.

Step 1: Navigate to **Managed Network > Configuration > WLANs**, and then click **+**.

Step 2: In the **New WLAN** window, on the **General** page, enter the following information, and then click **Next**.

- Name (ssid)—**Example-Employee**
- Primary usage—**Employee**
- Broadcast on—**All APs**
- Forwarding mode—**Tunnel**

Step 3: In the **New WLAN** window on the **VLANs** page, enter the following information, and then click **Next**.

- VLAN—**1** (this will be changed in a subsequent procedure)

7.3 Configure the Employee SSID Security

For the Wi-Fi network, you can use a WPA2-Personal passphrase, or you can choose to have every employee authenticate with a username and password using WPA2-Enterprise.

Note WPA2-Enterprise is used to enable 802.1X authentication for a wireless network. The wireless client can also authenticate against the RADIUS server using an EAP-TLS exchange, and the mobility controller 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 WPA2-Personal with passphrase access, choose option 1. If you are planning to use WPA-2 enterprise authentication, choose option 2.

Option 1: WPA2-Personal with Passphrase Access

Step 1: In the **New WLAN** window on the **Security** page, enter the following information, and then click **Next**.

- Security Level slider—**Personal**
- Key management—**WPA2-Personal**
- Passphrase—**[password]**
- Retype—**[password]**

The screenshot shows the 'New WLAN' configuration window with the 'Security' tab selected. A progress bar at the top indicates the current step. On the left, a security level slider is set to 'Personal'. The main configuration area includes:

- Key management:** A dropdown menu set to 'WPA2-Personal'. Below it, the radio button for 'Use static Pre-Shared Key (PSK)' is selected. The 'Passphrase' and 'Retype' fields contain masked text (dots).
- Auth servers:** A large empty box with a blue plus sign at the bottom right, indicating no servers are currently configured.
- MAC authentication:** A dropdown menu set to 'Disabled'.
- Blacklisting:** A toggle switch that is currently turned off.

Step 2: Skip to the next procedure.

Option 2: WPA2-Enterprise with Username and Password

Step 1: In the **New WLAN** window on the **Security** page, enter the following information, and then from inside the **Auth servers** box, click **+**.

- Security Level slider—**Enterprise**
- Key management—**WPA2-Enterprise**

The screenshot shows the 'New WLAN' configuration window with the 'Security' tab selected. A progress bar at the top indicates the current step. On the left, a security level slider is set to 'Enterprise', with 'More Secure' at the top and 'Less Secure' at the bottom. The main configuration area includes: 'Key management' set to 'WPA2-Enterprise'; 'Auth servers' with a '+' button; 'Reauth interval' set to '1440 min.'; 'Machine authentication' set to 'Disabled'; and 'Blacklisting' which is disabled.

Step 2: In the **Add Existing Server** window, select **Example-CPPM**, and then click **OK** and **Next**.

7.4 Configure the Employee SSID Access Rules

Step 1: In the **New WLAN** window on the **Access** page, enter the following information, and then click **Finish**.

- Default role—**authenticated**

New WLAN

General VLANs Security Access

Default role:

Server-derived roles:

Show [roles](#)

Step 2: Click **Pending Changes > Deploy changes > Close**.

7.5 Configure WLAN for Guest SSID

Step 1: Navigate to **Managed Network > Configuration > WLANs**, and then click **+**.

Step 2: In **New WLAN** window on the **General** page, enter the following information, and then click **Next**.

- Name (ssid)—**Example-Guest**
- Primary usage—**Guest**
- Broadcast on—**All APs**
- Forwarding mode—**Tunnel**

New WLAN

General VLANs Security Access

Name (SSID):

Primary usage: Employee Guest

Broadcast on:

Forwarding mode:

Step 3: In the **New WLAN** window on the **VLANs** page, enter the following information, and then click **Next**.

- VLAN—1 (this will be changed in a subsequent procedure)

7.6 Configure the Guest SSID Security

You can use WPA2-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.

Step 1: In the **New WLAN** window on the **Security** page, enter the following information, and then click **Next**.

- Security Level slider—**Internal captive portal with email registration**

Step 2: In the **New WLAN** window on the **Access** page, click **Finish**.

Note A new default role for this WLAN is created using the SSID name prepended to the guest-logon role name. The rules from the guest-logon role are copied to the new role. This allows you to customize the guest role for this WLAN without modifying the original guest logon.

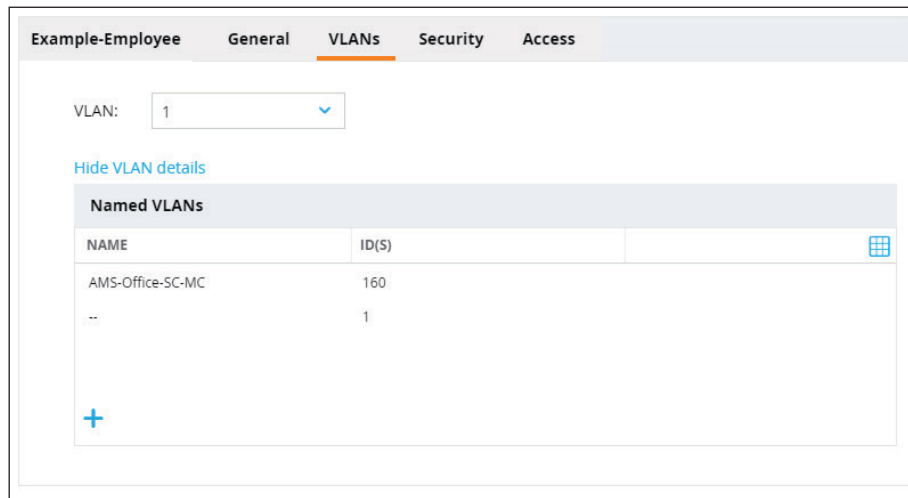
The default role for this guest WLAN: Example-Guest-guest-logon



Step 3: Click **Pending Changes > Deploy changes > Close**.

7.7 Configure Local VLANs for Employee and Guest SSIDs

Step 1: For the employee VLAN, navigate to **Managed Network > AMS > Office > Santa-Clara > Configuration > WLANs > Example-Employee > VLANs**, click **Show VLAN details**, and then click **+**.




Step 2: In the **New VLAN** window, enter the following information, and then click **Submit**.

- VLAN name—**AMS-Office-SC-Employee**
- VLAN ID/Range—**330**

Step 3: In the **Example-Employee** window, in the VLAN menu, select the name you created in the previous step, and then click **Submit**.

- VLAN—**AMS-Office-SC-Employee**



The screenshot shows the configuration page for the 'Example-Employee' WLAN. The 'VLANs' tab is active, and the 'VLAN' dropdown menu is set to 'AMS-Office-SC-Employee'. Below this, there is a section for 'Named VLANs' with the following table:

NAME	ID(S)
AMS-Office-SC-MC	160
AMS-Office-SC-Employee	330
--	1

Step 4: For the guest VLAN, navigate to **Managed Network > AMS > Office > Santa-Clara > Configuration > WLANs > Example-Guest > VLANs**, click **Show VLAN details**, and then click **+**.

Step 5: In the **New VLAN** window, enter the following information, and then click **Submit**.

- VLAN name—**AMS-Office-SC-Guest**
- VLAN ID/Range—**340**

Step 6: In the **Example-Guest** window, on the VLAN menu, select the name you created in the previous step, and then click **Submit**.

- VLAN—**AMS-Office-SC-Guest**

The screenshot shows the configuration page for a WLAN named 'Example-Guest'. The 'VLANs' tab is active, and the 'VLAN' dropdown menu is set to 'AMS-Office-SC-Guest'. Below this, there is a section titled 'Named VLANs' which contains a table of existing VLANs.


NAME	ID(S)
AMS-Office-SC-MC	160
AMS-Office-SC-Employee	330
--	1
AMS-Office-SC-Guest	340

Step 7: Click **Pending Changes > Deploy changes > Close**.

7.8 Configure Mobility Controller Interface with VLANs

Next, you configure the uplink interface in the mobility controller to allow user and guest VLANs to access the port.

Step 1: Navigate to **Managed Network > AMS > Office > Santa-Clara > CLS-MC-1 > Configuration > Interfaces > Ports**, and then, if using port channels, click **PC-0**.



The screenshot displays a network configuration interface with two main sections: "Port Channel" and "Ports".

Port Channel

NAME	MEMBERS	PROTOCOL	TRUSTED	POLICY	MODE	NATIVE VL...	TRUNK VL
PC-0	--	LACP	✓	Not-defined	trunk	160	160

Ports

PORT	ADMIN S...	TRU...	POLI...	MODE	NATI...	ACC...	TRU...	SPA...	MO...	DES...
GE-0/0/0	Enabled	--	Not...	access	1	1	1-40...	✓	--	--
GE-0/0/1	Enabled	--	Not...	access	1	1	1-40...	✓	--	--
GE-0/0/2	Enabled	--	Not...	access	1	1	1-40...	✓	--	--
GE-0/0/3	Enabled	--	Not...	access	1	1	1-40...	✓	--	--
GE-0/0/4	Enabled	✓	Not...	PC-0	160	1	1-40...	✓	--	GE0/...

Step 2: In the **PC-0** window, enter the following information, and then in the **Allowed VLANs** section, click **+**.

- Mode—**Trunk**
- Allowed VLANs—**Allow specified VLANs**

The screenshot displays the configuration page for a Port Channel named PC-0. The top navigation bar includes 'Ports', 'VLANs', 'IP Routes', 'IPv6 Neighbors', and 'GRE Tunnels'. The 'Port Channel' section contains a table with the following data:

NAME	MEMBERS	PROTOCOL	TRUSTED	POLICY	MODE	NATIVE VL...	TRUNK V
PC-0	--	LACP	✓	Not-defin...	trunk	160	160

Below the table is a '+' icon. The 'PC-0' configuration details are as follows:

- Port channel ID: PC-0
- Protocol: LACP
- LACP mode: active
- Port members: GE-0/0/4, GE-0/0/5 (with an 'Edit' button)
- Admin state:
- Trust:
- Policy: Not-defined
- Mode: Trunk
- Native VLAN: 160
- Allowed VLANs: Allow specified VLANs

The 'Allowed VLANs' section contains a table with the following data:

VLAN	TRUSTED / ...
160	Trusted

A '+' icon is located at the bottom of the 'Allowed VLANs' section.

Step 3: In the **Add Allowed VLAN** window, enter the following information, and then click **OK** and **Submit**.

- VLAN—**160,330,340**

Step 4: For each additional mobility controller, repeat this procedure, changing the variables as required.

Step 5: Click **Pending Changes > Deploy changes > Close**.


7.9 Configure Mobility Controller Guest VLAN with an IP Address

A guest captive portal is a method of authentication that presents a web page that requires action on the part of the user before network access is granted. The required action can be agreeing to an acceptable-use policy or entering a user ID and password, which are validated against a database of authorized users. The captive portal requires an IP address in order to reply to the client and send the web portal redirection to the client.

Configure a DHCP IP address on the guest VLAN of the mobility controller.

Step 1: Navigate to **Managed Network > AMS > Office > Santa-Clara > CLS1-MC-1 > Configuration > Interfaces > VLANs**, click **AMS-Offices-SC-Guest**, on the **VLANs** section click **340**, on **Port Members** section below, click **IPv4**, enter the following information, and then click **Submit**.

- IP assignment—**DHCP**



The screenshot shows a configuration page with tabs for 'Port Members', 'IPv4', 'IPv6', and 'More'. The 'IPv4' tab is selected. Under the 'IP Address Assignment' section, there are three fields: 'IP assignment:' with a dropdown menu set to 'DHCP', 'Client ID:' with an empty text box, and 'MTU:' with an empty text box. Below these fields are two expandable sections: 'IGMP' and 'Other Option', both with a right-pointing chevron.

Repeat this step for the guest VLANs on each mobility controller.

Step 2: Click **Pending Changes > Deploy changes > Close**.

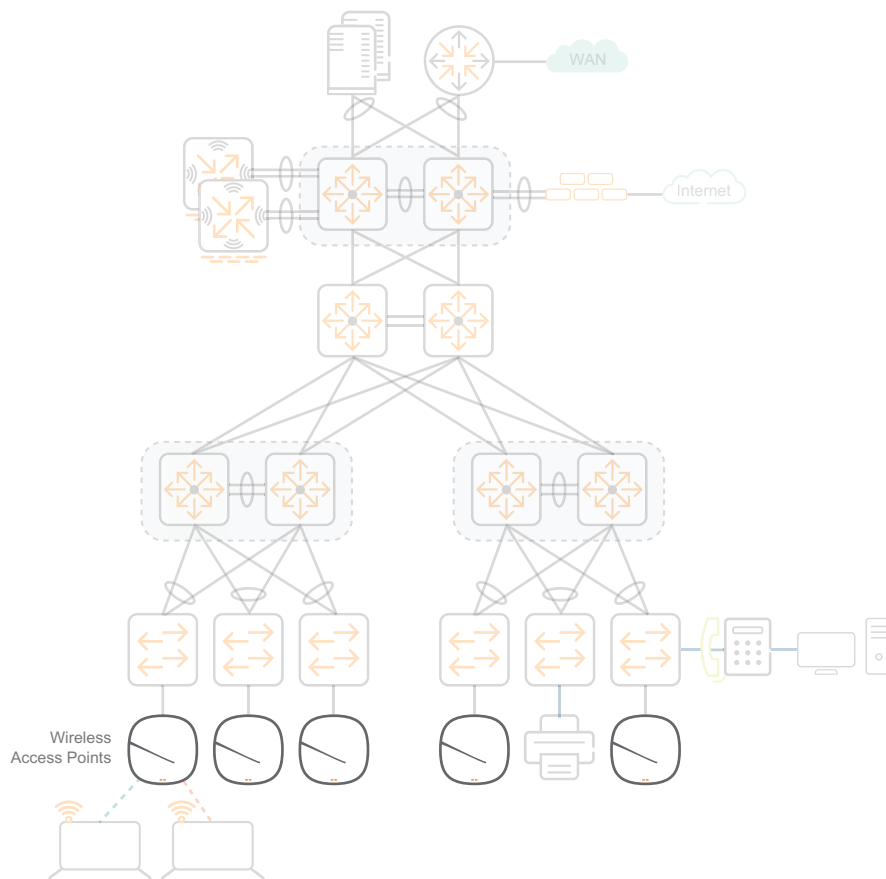
Procedures

Configuring the Access Points

- 8.1 Configure AP Group
- 8.2 Remove SSID from Default AP Group
- 8.3 Configure the Power Settings
- 8.4 Disable 80 MHz-wide Channels in the 5 GHz Band
- 8.5 Enable Control Plane Security and Auto Certificate Provisioning
- 8.6 Convert Instant Access Points to a Campus Access Points
- 8.7 Provision Access Points to the Mobility Controller AP Group

Use this section to configure the access points. The following figure shows the wireless access points in the access layer of the Aruba Campus design.

Figure 37 Aruba Campus design—wireless access points



1068A

8.1 Configure AP Group

Use an AP group to apply a set of features to a group of APs. You can also configure a feature for a specific AP. Any values that you configure for a specific AP override the same values configured for the AP group to which the AP belongs.

Step 1: Navigate to **Managed Network > Configuration > AP Groups**, and then click **+**.

Step 2: In the **New AP Group** window, enter the following information, and then click **Submit**.

- Name—**Example-AP-Group**

Step 3: Navigate to **Managed Network > Configuration > AP Groups > Example-AP-Group > WLANs**, and then click **+**.

Step 4: In the **Select WLAN** window, enter the following information, and then click **Submit**.

- Virtual-ap—**Example-Employee**

Step 5: For the guest WLAN, repeat Step 3 and Step 4, changing the variable as required.

Step 6: Click **Pending Changes > Deploy changes > Close**.

8.2 Remove SSID from Default AP Group

Next, you remove the employee and guest SSIDs from the default AP group. An *AP group* is a set of APs to which the same configuration is applied. There is an AP group called “default” to which all APs are assigned.

Step 1: Navigate to **Managed Network > Configuration > AP Groups > default > WLANs**, select **Example-Employee**, and then click the **trash can** icon on the right.

Step 2: In the **Are you sure you want to delete** window, click **Delete**.

Step 3: For the guest WLAN, repeat Step 1 and Step 2.

Step 4: Click **Pending Changes > Deploy changes > Close**.

8.3 Configure the Power Settings

Next, you configure the power settings for the 5-GHz and 2.4-GHz radios. In the 2.4-GHz band, set the 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. In the 5-GHz band 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 continue to provide normal access service to clients. They 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 **Managed Network > Configuration > AP Groups > Example-AP-Group > Radio**.

Step 2: In the **2.4 GHz** section, enter the following information.

- Radio mode—**ap-mode** (default)
- Spectrum monitoring—**Enabled**
- Min Transmit EIRP (dBm) slider—**6**
- Max Transmit EIRP (dBm) slider—**9**

Step 3: In the **5 GHz** section, enter the following information, and then click **Submit**.

- Radio mode—**ap-mode** (default)
- Spectrum monitoring—**Enabled**
- Min Transmit EIRP (dBm) slider—**12**
- Max Transmit EIRP (dBm) slider—**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 difference between the min and max Transmit EIRP settings should not exceed 6 dBm for all radios.



AP Groups > Example-AP-Group APs WLANs **Radio** Mesh LMS MultiZone

▼ **Basic**

2.4 GHz	5 GHz
Radio mode: ap-mode	Radio mode: ap-mode
Spectrum monitoring: <input checked="" type="checkbox"/>	Spectrum monitoring: <input checked="" type="checkbox"/>
Transmit EIRP(dBm): 6 - 9	Transmit EIRP(dBm): 12 - 15
Valid channels: 1,6,11 1-5,7-11 Edit	Valid channels: 36,40,44,48,149,153,157,161,165 36-40,44-48,149-153,157-161 36-48,149-161 Edit

> Advanced
> Client Control

Step 4: Skip to the next procedure.

Option 2: Walled-office Environment

Step 1: Navigate to **Managed Network > Configuration > AP Groups > Example-AP-Group > Radio**.

Step 2: In the **2.4 GHz** section, enter the following information.

- Radio mode—**ap-mode** (default)
- Background spectrum monitoring—**Enabled**
- Min Transmit EIRP (dBm) slider—**6**
- Max Transmit EIRP (dBm) slider—**9**

Step 3: In the **5 GHz** section, enter the following information, and then click **Submit**.

- Radio mode—**ap-mode** (default)
- Spectrum monitoring—**Enabled**
- Min Transmit EIRP (dBm) slider—**15**
- Max Transmit EIRP (dBm) slider—**18**

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 difference between the min and max Transmit EIRP settings should not exceed 6 dBm for all radios.



AP Groups > Example-AP-Group APs WLANs **Radio** Mesh LMS MultiZone

▼ **Basic**

2.4 GHz	5 GHz
Radio mode: ap-mode	Radio mode: ap-mode
Spectrum monitoring: <input checked="" type="checkbox"/>	Spectrum monitoring: <input checked="" type="checkbox"/>
Transmit EIRP(dBm): 6 - 9	Transmit EIRP(dBm): 15 - 18
Valid channels: 1,6,11 1-5,7-11	Valid channels: 36,40,44,48,149,153,157,161,165 36-40,44-48,149-153,157-161 36-48,149-161

▶ Advanced


▶ Client Control

8.4 Disable 80 MHz-wide Channels in the 5 GHz Band

Step 1: At the bottom of the **Radio** tab in the **5 GHz** section, click **Edit**.


Step 2: On the **Valid 5 GHz Channels** page in the **UNII-1** section, disable the 80 MHz channels by matching the colors as depicted below.

- 80 MHz—**Gray** (disabled)
- 40 MHz—**Orange** (enabled)
- 20 MHz—**Orange** (enabled)



Repeat this step for the remaining UNII sections on the page.

Step 3: After all sections are complete, click **OK**, and then click **Submit**.



Step 4: Click **Pending Changes > Deploy changes > Close**.

8.5 Enable Control Plane Security and Auto Certificate Provisioning

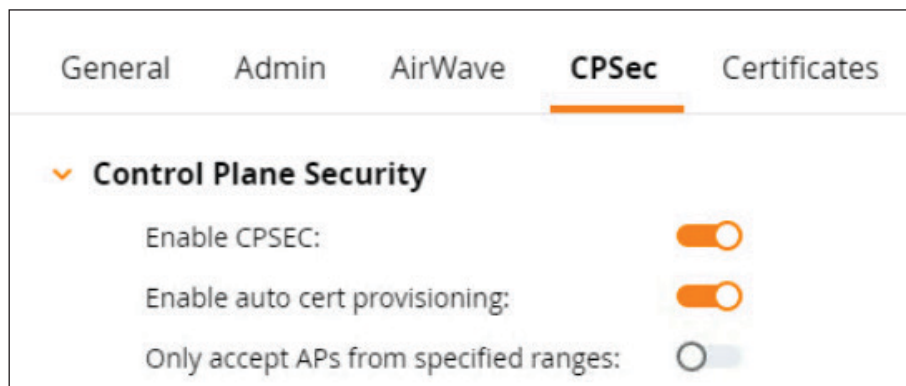
Next, you enable Control Plane Security (CPSec) and auto certificate provisioning. ArubaOS supports secure IPsec communications between a controller and campus APs using public-key self-signed certificates created by each Mobility Master. The controller certifies its APs by issuing them certificates. If the Mobility Master has any associated mobility controllers, it sends a certificate to each mobility controller, which in turn sends certificates to their associated APs.

The mobility controller maintains a whitelist that contains records of all campus APs connected to the network. The campus AP whitelist is used to add valid campus APs to the secure network or revoke network access to any suspected rogue or unauthorized APs. If all APs on network are valid, such as during the initial configuration, automatic certificate provisioning can be enabled to send certificates from the controller to each AP. This automatically provisions them and adds them to the AP whitelist.

When the controller sends a certificate to the AP, that AP must reboot before it can connect to the controller over a secure channel. If you are enabling CPsec for the first time on a large network, your users may experience several minutes of interrupted connectivity while each AP receives its certificate and establishes its secure connection.

Step 1: Navigate to **Managed Network > Configuration > System > CPsec > Control Plane Security**, slide the buttons to the right for the selections below, and then click **Submit**.

- Enable CPSEC slider—**Enable**
- Enable auto cert provisioning slider—**Enable**



Step 2: Click **Pending Changes > Deploy changes > Close**.

8.6 Convert Instant Access Points to a Campus Access Points

Optional

This optional procedure converts instant access points (IAPs) to campus access points (CAPs). If you have existing IAPs or APs running older software and you want the Aruba mobility controller to manage them, they need to be converted to CAP mode using the following steps.

If you have newer APs running the latest software, you can skip this procedure by allowing the AP to discover the controllers by using the information from 8.7.

Step 1: Connect the AP 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 automatically opens 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 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 AP.



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

- Username—**admin**
- Password—**admin** (default)

Step 5: At the top of the page on the right hand side, navigate to **Maintenance > Convert**, enter the following information, and then click **Convert Now**.

- Convert one or more Access Points to—**Campus APs managed by a Mobility Controller**
- Hostname or IP Address of Mobility Controller—**10.2.160.100** (VRRP IP address)

Step 6: In the **Confirm Access Point Conversion** window, click **Convert Now**. The conversion process takes several minutes, and then the APs reboot to join the cluster.

8.7 Provision Access Points to the Mobility Controller AP Group

An AP can discover the IP address of the controller from a DNS server, from a DHCP server, or using the Aruba Discovery Protocol.

At boot time, the AP builds a list of managed device IP addresses and then tries these addresses in order until it successfully reaches a managed device. This list of IP addresses provides an enhanced redundancy scheme for managed devices that are located in multiple data centers separated across layer-3 networks.

The AP constructs its list of managed device addresses as follows:

- If the provisioning parameter is set to a DNS name, that name is resolved and all resulting addresses are put on the list. If is set to an IP address, that address is put on the list.
- If the provisioning parameter is not set and a managed device address was received in DHCP option 43, that address is put on the list.
- If the provisioning parameter is not set and no address was received via DHCP option 43, you use AP discovery protocol (ADP) to discover a managed device address and that address is put on the list.
- Managed device addresses derived from the server-name and server-ip provisioning parameters and the default managed device name aruba-master are added to the list. Note that if a DNS name resolves to multiple addresses, all addresses are added to the list.

Most DHCP servers can send a variety of optional information, including the Vendor-Specific Option Code, also called option 43. Here is how option 43 works for an Aruba AP:

1. The DHCP client on an Aruba AP adds an optional piece of information called the Vendor Class Identifier Code (Microsoft DHCP option 60) in its DHCP request. The value of this code is: **ArubaAP**
2. The DHCP server sees the Vendor Class Identifier Code in the request and checks to see if it has option 43 configured. If it does, it sends the Vendor-Specific Option Code (option 43) to the client. The value of this option is the VRRP IP address of the Aruba mobility controller cluster.
3. The AP receives a response from the DHCP server and checks if option 43 is returned. If it is, the AP contacts the mobility controller cluster using the supplied IP address.

Microsoft Windows-based DHCP Server

The Microsoft Windows-based DHCP server requires option 60 and option 43 configured for an Aruba AP. Because option 60 is not a predefined option on a Windows DHCP server, you must add it to the option list for the server. The following information is required when you are creating option 60:

- Name—**Aruba access point**
- Data—**Type String**
- Code—**60**
- Description—**Aruba AP vendor class identifier**

When option 60 is added to the DHCP scope, the following string value is required:

- String value—**ArubaAP**

When option 43 is added to the DHCP scope, the following value is required:

- ASCII—**10.2.160.100** (VRRP IP address of the mobility controller cluster)

After the APs have discovered the controllers in your network, they need to be provisioned in an AP group.

Step 1: Navigate to **Managed Network > AMS > Office > Santa-Clara > Configuration > Access Points > Campus APs**, select the AP and then click **Provision**.

Step 2: In the window with AP's MAC address below the Campus APs window, enter the following information, and then click **Submit**.

- Name—**515-AG2-AC3-11**
- AP group—**Example-AP-Group**
- Controller discovery—**Static**
- Controller IP/DNS name—**10.2.160.100** (VRRP IP address of mobility controller cluster)
- IP—**DHCP**
- Deployment—**Campus**

The screenshot shows a configuration window for an AP with the title "515-AG2-AC3-11". The window contains the following fields and options:

- MAC address: 9c:8c:d8:c9:10:26
- Name: 515-AG2-AC3-11
- AP group: Example-AP-Group (dropdown menu)
- Controller discovery: Use AP discovery protocol (ADP) Static
- Controller IP/DNS name: 10.2.160.100
- IP: DHCP Static
- Deployment: Campus Remote Mesh Remote mesh portal
- Wifi uplink:

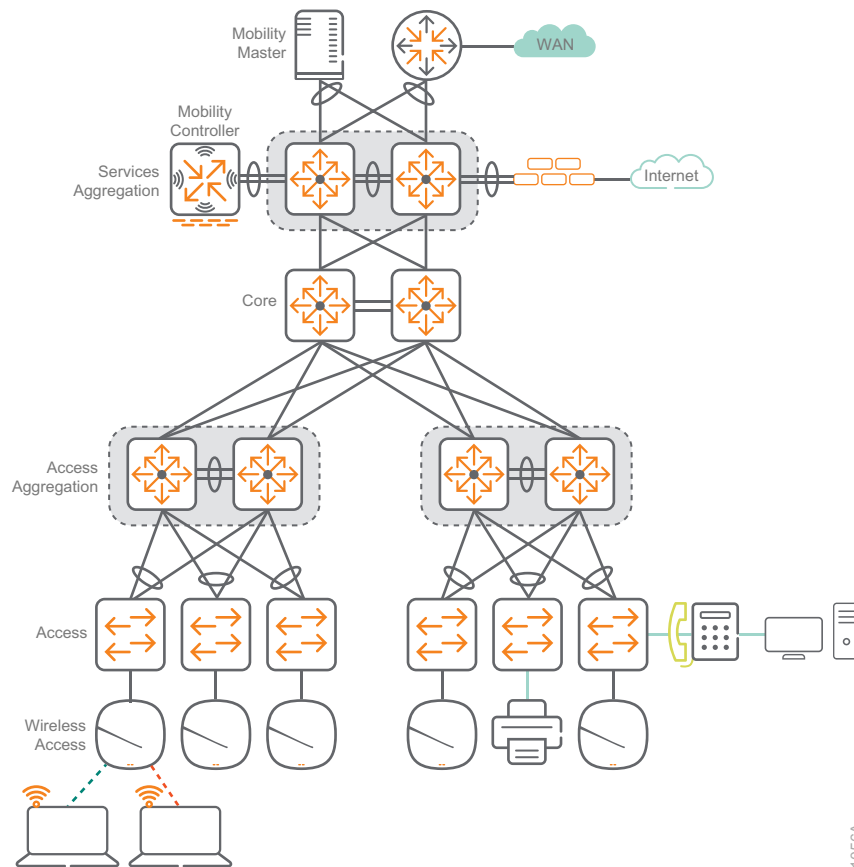
Step 3: On the **Access Points will be Rebooted** screen, click **Continue & Reboot**.

Step 4: For each additional AP you want to provision, repeat Step 1 through Step 3.

Summary

The flow of information is a critical component to a well-run organization. The Aruba 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, which helps improve user satisfaction and productivity while reducing operational expense.

Figure 38 Aruba Campus design



The Aruba 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 hardware and software were validated for this guide:

Wired Core

Product name	Software version
Aruba 8400	10.03.0040

Wired Aggregation

Product name	Software version
Aruba 8320	10.03.0040
Aruba 5400R	16.08.006
Aruba 3810M	16.08.006

Wired Access

Product name	Software version
Aruba 5400R	16.08.006
Aruba 3810M	16.08.006
Aruba 2930M	16.08.006
Aruba 32930F	16.08.006

Wireless Mobility Master

Product name	Software version
MM-VA-1K	8.5.0.3 (ArubaOS)

Wireless Mobility Controller

Product name	Software version
Aruba 7205	8.5.0.3 (ArubaOS)

Wireless Access Points

Product name	Software version
Aruba 515 Series AP	8.5.0.3 (ArubaOS)
Aruba 340 Series AP	8.5.0.3 (ArubaOS)
Aruba 330 Series AP	8.5.0.3 (ArubaOS)
Aruba 300 Series AP	8.5.0.3 (ArubaOS)

What's New in This Version

We made the following changes since Aruba last published this guide:

- Updated information on Wi-Fi 6 and 5xx series APs
- Updated information on new Wi-Fi security features
- New VSX design for AOS-CX
- Updated access security recommendations for loop prevention
- Revised scale recommendations
- Updated software across all platforms
- Inclusion of Aruba 515 Series APs in validation
- Inclusion of 7205 Mobility Controllers in validation



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