aruba
a Hewlett Packard Enterprise company

ESP Campus Design

Validated Solution Guide

Aruba Solution TME

June 2021

Table of Contents

ESP Campus Design

This guide is intended to assist an IT professional understand the following design considerations for a campus environment:

- Hardware selection
- Software selection
- Topology
- High availability
- Scalability
- Application performance
- Security

NOTE:

For the most up-to-date information on ESP Campus solutions, please visit the following: Validated Solution Guide Program

Introduction

Like previous technology transitions, shifting to the “era of data” at the Intelligent Edge changes the role of the network infrastructure and introduces new challenges. Corporate networks have always played a pivotal role in moving data and connecting people to their apps and services. With the Intelligent Edge, the network requirements go far beyond standard connectivity and the antiquated access technologies of the past.

The need to support robust networking and new application services places incredible demands on IT staff. The challenges are significant given the existence of legacy infrastructure that is inherently siloed, requires excessive manual network operations, and is complex and difficult to manage. Network issues can only be resolved when they are found and addressed by humans. Often, troubleshooting network problems is like finding needles in a haystack. IT leaders need to carefully assess their infrastructure and operational models to ensure the network, tools and operator experience are poised to support business success in this new era using a modern architecture.

Purpose of This Guide

This design guide covers the Campus in the Edge Services Platform (ESP) architecture, 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 they will provide to an organization. The guide describes a single system that integrates access points, gateways, access switches, aggregation switches, core switches, and cloud-based orchestration and network management.

Design Goals

The overall goal is to create a simple scalable design that is easy to replicate at different sites. 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. 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 recommended designs, features, software and hardware.

Audience

This guide is written for IT professionals who need to design Aruba solutions for small, medium and large networks. These IT professionals can fill a variety of roles:

- Systems engineers who need a standard set of procedures for implementing Aruba 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.

The Wi-Fi 5 and Wi-Fi 6 standards support speeds greater than 1 Gbps. To accommodate the increased data rates, the APs implement the IEEE 802.3bz Ethernet standard of 2.5 and 5 Gbps. An organization can achieve the higher data rates on existing building twisted-pair cabling when connecting to Aruba switches with Smart Rate ports which also support the 802.3bz Ethernet standard. 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 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 MAC authentication and profiling to prevent rouge 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.

This guide discusses the following use cases:

- Artificial Intelligence to augment the operator with Smart Telemetry
- Zero Trust Security to secure the network from inside and outside attacks
- Unified Infrastructure with centralized cloud-based management

Aruba ESP Architecture

The Aruba Edge Services Platform (ESP) is an evolution of Aruba's end-to-end architecture, providing a Unified Infrastructure with centralized management leveraging Artificial Intelligence Operations (AIOps) for improved operational experience that helps enable a Zero Trust security policy on an existing infrastructure. Aruba ESP is the industry's first platform that is purpose-built for the new requirements of the Intelligent Edge.

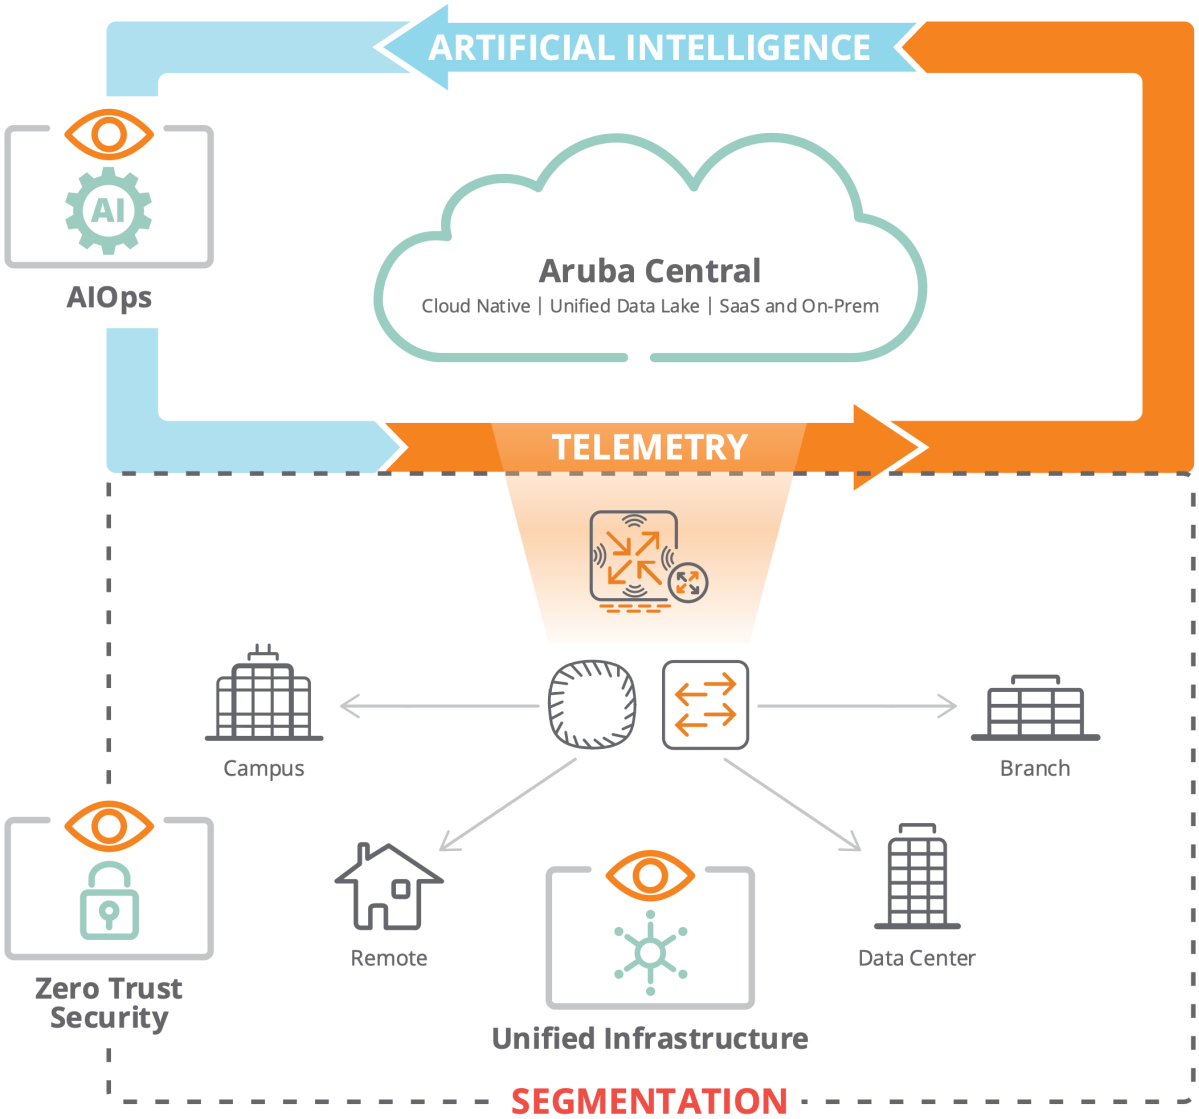

Figure 1: ESP Architecture

AI Operations

AIOps is not a product but rather a capability within the Aruba ESP architecture. This broad industry term comes from the combination of AI and IT operations. It refers to the way an IT team manages data and information. The ESP AIOps platform collects large amounts of data from the network and presents it in a meaningful way to the administrator. Aruba ESP uses Machine Learning (ML) and analytics capabilities to reveal network issues before users notice them.

By definition, ML is an application of Artificial Intelligence (AI) that provides systems the ability to automatically learn and improve from experience without being explicitly programmed. The primary goal of ML is to allow computers to learn automatically without human intervention or assistance and adjust actions accordingly. Aruba's use of ML helps the IT professional move faster to meet the service-level expectations of users and devices on the network.

Central sits at the core of Aruba's AIOps capability for wired, wireless and SD-WAN infrastructure. Designed using a modern microservices model, Central provides operations teams with a comprehensive view of application, user, and device status. It also provides network health, bandwidth usage, predictive insights, and prescriptive guidance through a single pane of glass. Data-driven dashboards allow the network administrator to monitor, identify, and address issues before they impact the business, and peer benchmarking shows them how to adjust for optimal performance. The operations team can easily view network insights across any location, from one central place, because the data is stored securely in the cloud.

Unified Infrastructure

With the network serving as the center of convergence, new expectations of performance, density and reliability leave enterprises with the need to optimize their infrastructure for the cloud. Configuration changes on a per port basis, hardware sizing, and even day-to-day maintenance can be automated and delivered as a service with minimal overhead or disruption for both IT and end-users alike.

To allow flexible deployment options, the Aruba ESP infrastructure can be implemented in both physical and virtual form factors. Aruba has physical switches, Gateways and APs, as well as SD-Branch virtual Gateways. The ESP architecture components can also be deployed on-premises or in the cloud. By supporting a variety of clouds, an organization can connect and secure physical locations, private clouds, or public clouds in a consistent manner. Using a common data lake, Central correlates and displays multiple dimensions of information in context. It also unlocks powerful capabilities around automated root cause analysis, predicts issues before they impact the business and provides robust analytics. By automating simple tasks across the lifecycle of the network, administrators can focus on driving innovation and using the network to create business value at the Edge.

Zero Trust Security

One of the major capabilities of Aruba ESP is the ability to enable a Zero Trust Security policy. Zero Trust uses a different approach than traditional methods of network security. First, it assumes there are attackers both inside and outside of the network. Zero Trust starts with a default "deny all" posture, which means no user or device seeking access to the network is trusted. A strict identity verification is required of all users and devices, regardless of whether they are inside or outside the known network boundaries.

The second aspect of Zero Trust is access control. The policy grants access based on the identity and role of the user, the device from which they are connecting, and their connection context, such as date and time, geographic location, and a security posture score. Zero Trust adaptively provides the appropriate access required in a "least-privilege access" response. This means users receive only as much access to resources as they need at the time of request, and further access is only granted when absolutely required. This greatly minimizes the exposure to sensitive parts of the network.

The final element of Zero Trust is continuous monitoring and assessment. If a user and their device are potentially compromised, then the combined access permissions are immediately modified to reflect the increased risk from the user and the associated device. This means a user accessing information from their normal location on their corporate desktop computer in their physical office is not denied access when someone has fraudulently used their credentials to access network resources from an outside network because a Zero Trust policy takes into account the user, device and location information.

Aruba ESP Architecture Layers

Aruba ESP offers a breadth of services, including onboarding, provisioning, orchestration, analytics, location tracking and management. AI Insights reveal issues before they impact users allowing an organization to accomplish tasks quickly and easily with intuitive workflow-centric navigation using views that present multiple dimensions of correlated data. Policies are created centrally and features like Dynamic Segmentation allow the network administrator to implement them over an existing infrastructure. This is possible because the Aruba ESP architecture is built in distinct layers, as shown in the following figure.

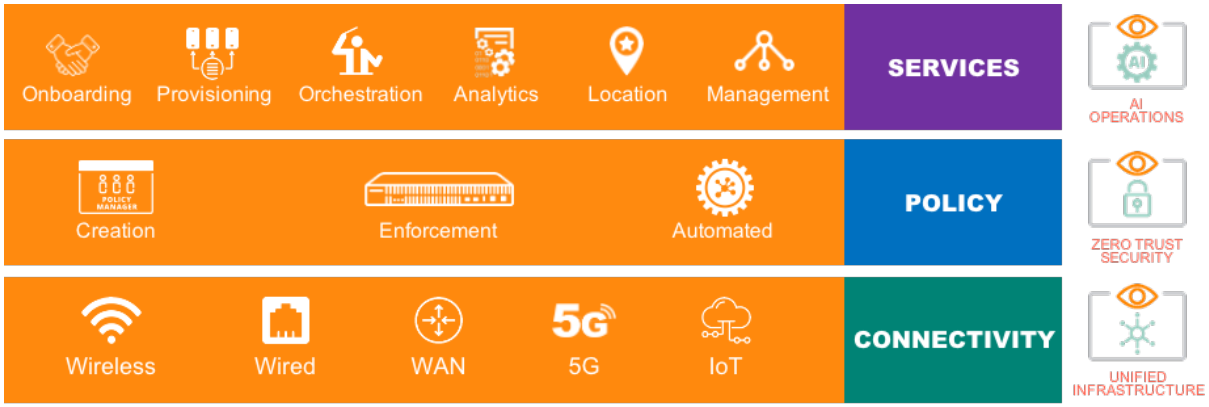

Figure 2: ESP layers

Connectivity Layer

Starting at the bottom, the *Connectivity Layer* is the foundation for the Aruba ESP architecture. It features the physical infrastructure components with extensive flexibility and high availability, along with the telemetry for analyzing the current state of the network. Enterprise networks need to support their existing endpoints, including many legacy systems, while transitioning to a modern architecture that is optimized for the new use cases at the Intelligent Edge.

Architectures that require a full network redesign, including hardware and network protocol stack, can cause serious interruption and risk-compatibility issues, especially with a company's legacy systems. With Aruba ESP, legacy devices have the option to continue operating in the connectivity layer, and enhancements to the policy layer running on top allow for additional security and control on the infrastructure beyond what is deployed in traditional networks. The components of the connectivity layer consist of the Gateways, APs with wireless and IoT radio capabilities, and switches as shown in the following figure.

Figure 3: Connectivity Layer

Policy Layer

The *Policy Layer* is responsible for the creation, maintenance and enforcement of network security policies. An endpoint security posture is determined from information gathered using context from the network. After a determination is made, a user role is assigned, and policy can be applied based on the trust level of the device. The ESP architecture supports consistent policy for wired and wireless that includes options for access control at the edge of the network and also traffic tunneling for additional inspection by the stateful firewall on the Gateway.

A user's context is easily shared to other domains and existing VLAN and IP addresses structures are maintained, but because policy is enforced at the user and group levels, VLANs and IP addresses are no longer tied to policy. Some of the additional features enabled by a flexible policy layer include traffic segmentation, service insertion and transport independence. Management of the policy layer is done with Central and ClearPass Policy Manager, while the enforcement is handled by the Gateway, switch or AP depending on the specifics of the design.

Figure 4: Policy Layer

Services Layer

The *Services Layer* is where the operations team interacts with the Connectivity and Policy layers. It provides significant capabilities leveraging AI, ML and location-based services for network visibility and insights into how the network is performing. By leveraging a unified data lake in the cloud, Aruba ESP correlates cross-domain events and displays multiple dimensions of information in context, unlocking powerful capabilities around automated root-cause analysis while providing robust analytics. The primary home for Services Layer functionality is Central. However, there are components which provide additional services like user assurance and location tracking.

Figure 5: Services Layer

Aruba offers multiple deployment models to meet the varying business and technical requirements of an organization. Most organizations purchase and deploy Aruba's portfolio of Gateways, switches and APs and manage them from the cloud. The small number of organizations with security policies preventing use of public cloud services can deploy the ESP services layer on-premises or in a private cloud. Lastly, organizations that prefer their network as-a-service can deploy it as a managed service through HPE Greenlake for Aruba.

ESP Campus Architecture

The ESP Campus architecture consists of wired and wireless options to accommodate different use cases based on physical environments, features, and scaling requirements. Innovations in high availability, upgrade simplicity, and programmability are what separates the Aruba ESP Campus from the competition.

Campus Wired Architecture

The Aruba CX portfolio provides a variety of form factors which include models with the latest networking standards. These switches are available in modular and stackable options to satisfy a diverse set of requirements. They use a cloud-native operating system called AOS-CX which is designed with a focus on network resiliency utilizing a database centric operational model. With features like always-on PoE, Virtual switching Framework (VSF), and Virtual switching Extension (VSX), organizations can rely on the network infrastructure for their mission critical traffic.

AOS-CX allows a variety of two-tier and three-tier options from a Layer 3 routed access layer all the way to switched Layer 2 access with a redundant Layer 3 aggregation and core. Most large organizations adopt a model where the aggregation layer provides Layer 2 towards the access switches and Layer 3 towards a fully routed core. To ensure this design has maximum resiliency without added complexity, Aruba created a high availability system that supports multi-chassis link-aggregation (MC-LAG) while keeping the control plane of each switch independent. This capability is called Aruba VSX and provides a redundant, loop-free topology that does not require the spanning tree protocol. VSX also provides DHCP redundancy, native active-active default gateway, and active-active Layer 2 and Layer 3 forwarding without blocked uplinks.

AOS-CX switches have multiple ways to identify devices and dynamically allow traffic upstream based on client authentication with a feature called a User Role. The first option is the Layer 2 bridged model which tags traffic with a VLAN and forwards it into the underlay. The next technique is User-Based Tunneling (UBT) which encapsulates traffic and sends it to an Aruba Gateway to have additional data plane services applied. UBT is the same as the tunnel mode used in the APs and provides a unification of policies that can be applied to traffic whether it is wired or wireless. Finally, Virtual Extensible LAN (VXLAN) encapsulates traffic and sends it to another switch which can be at the access edge or in the data center. These wired models give customers the tools to create flexible policies in the network while providing an option to continue to use traditional approaches such as VLANs.

The foundation of each reference architecture provided in this guide is the connectivity layer for the campus LAN. This design separates the network into smaller modular components. This is commonly referred to as a set of interconnected layers such as core, aggregation and access that form the main network along with additional services that provide specific functions such as Internet, WAN, Wireless and server aggregation.

This modular approach simplifies the overall design and management of the LAN while providing the following benefits:

- Modules can be easily replicated providing scale as the network grows.
- Modules can be added and removed with minimum impact to other layers as network requirements evolve.
- Modules allow the impact of operational changes to be constrained to a smaller subset of the network
- Modules provide specific fault domains improving resiliency.

NOTE:

The modular design philosophies outlined in this section are consistent with industry leading practices and can be applied to any size network.

Two-tier Wired

The two-tier wired architecture includes access switches or switch stacks connected to a dual-switch aggregation layer that also serves as a collapsed core. The access switches provide Layer 2 services to connected endpoints and connect up to aggregation switches providing both Layer 2 and Layer 3 services.

The two-tier modular design is well suited for small buildings with few wiring closets and access switches. It also works well in larger environments when the fiber cable runs from each wiring closet are homed into a single location. The following figure shows a two-tier wired model with a Layer 2 access, and Layer 3 collapsed core and services aggregation.

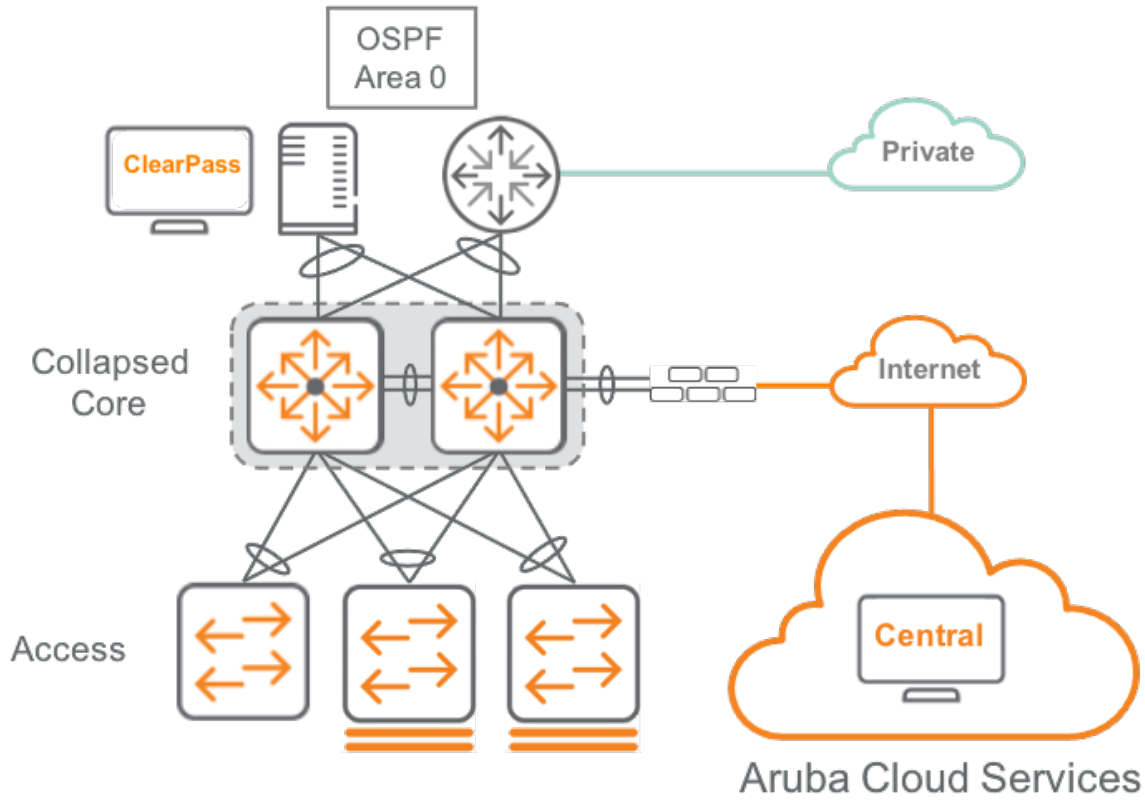

Figure 6: Two-tier wired with collapsed core and services

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 compared to a stack of fixed access switches. In the two-tier design, the collapsed core also provides critical network services like WAN aggregation, policy gateways, Internet DMZ and data center servers for a small to midsize organization.

Three-tier Wired

The three-tier wired architecture builds on the two-tier using the same access and aggregation layers in a repeatable fashion. An organization uses this design when the number of aggregation switches or the layout of the physical wiring plant does not work for the simpler two-tier model. When access layer switches are located in multiple geographically dispersed buildings, costly fiber-optic runs between buildings can be avoided 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. The decision to choose a two-tier or three-tier network design is based on overall network size or the physical layout of the wiring.

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, service aggregation layer switches for data center, WAN connectivity, and Internet edge services can be co-located. If the network is not large enough to warrant a standalone core, 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 the aggregation layer switch blocks.

In this design, the access switches are Layer 2, and the aggregation and core switches are Layer 3. The core switches must support 40 Gbps ports with enough capacity to grow and several 100 Gbps ports to connect between themselves. For large networks with thousands of users or where the physical wiring layout is not suited for a two-tier design, the three-tier campus design is recommended. The following figure shows a three-tier wired model with a Layer 2 access, and Layer 3 access aggregation, standalone core, and services aggregation.

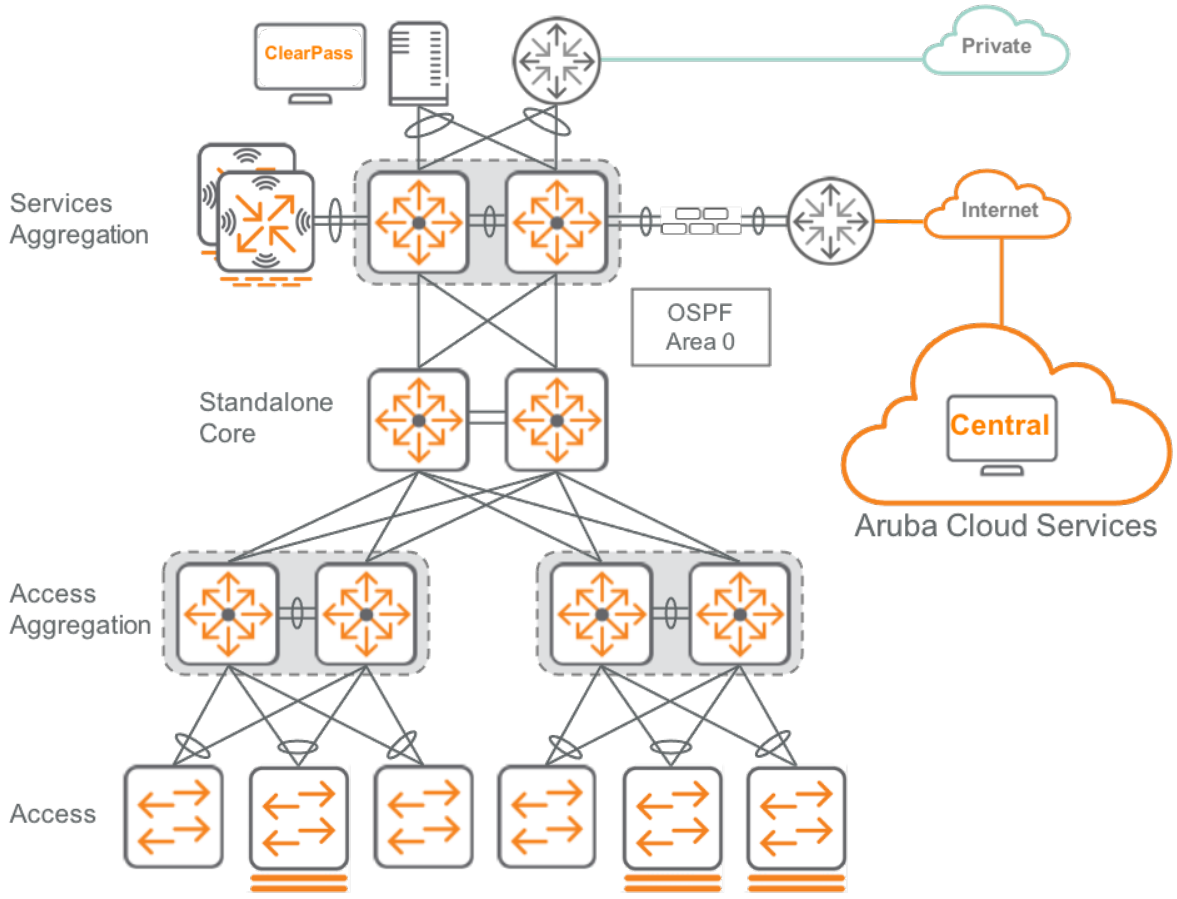

Figure 7: Three-tier wired with standalone core

The core and services can also be collapsed into a single pair of switches. When combining the core and services functions, stacking the two core-services switches together as a VSX pair is recommended 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 a standalone core is needed later, re-use the combined core-services switches as the services aggregation. The following figure shows a three-tier model with a Layer 2 access, and Layer 3 access aggregation and Layer 3 collapsed core and services aggregation.

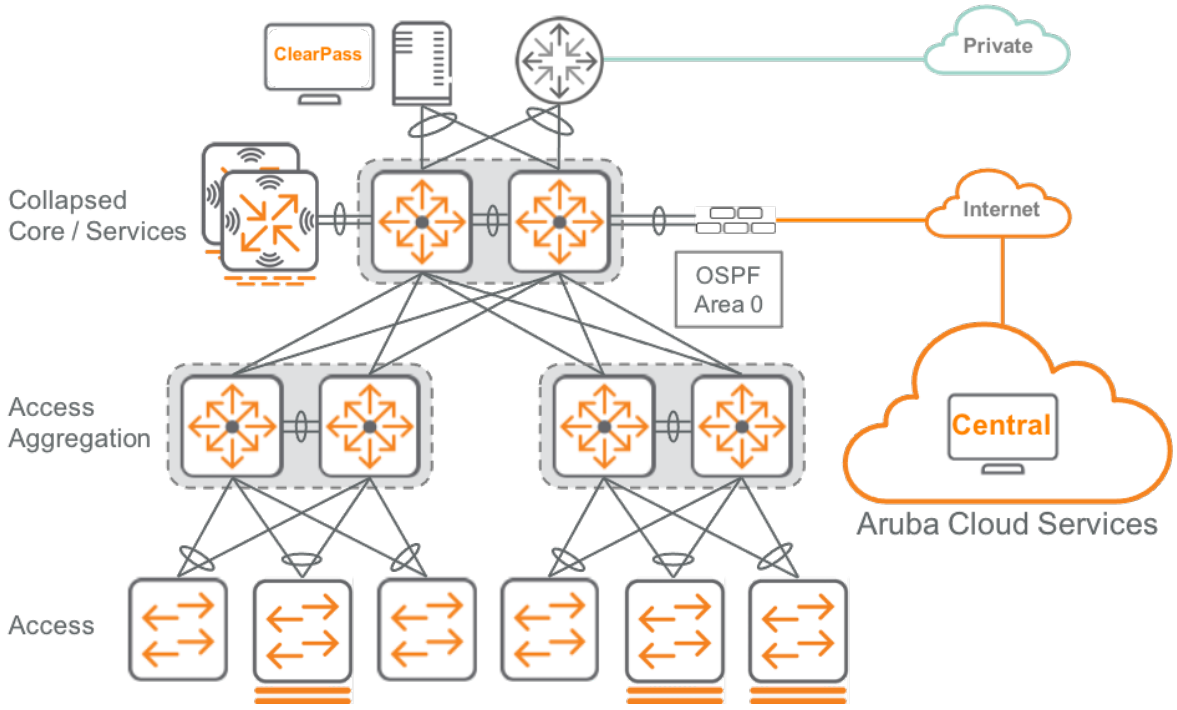

Figure 8: **Three-tier wired with collapsed core and services**

Access Layer

The access layer provides Layer 2 connectivity to the network for wired and wireless devices. As such, 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, making sure the devices cannot provide unauthorized services to end users, and prevents unauthorized devices from taking over the role of other devices on the network. The access layer provides automated services like PoE, QoS, and VLAN assignments in order to reduce operational requirements. To simplify the network as much as possible, the access layer is a Layer 2 design. Aruba aggregation switches can easily accommodate thousands of devices, so prior best practices of limiting broadcast domain sizes can be greatly expanded to accommodate larger subnets without impact on performance or usability.

Aggregation Layer

The aggregation layer's primary function is to give access switches a common connection point and to act as the boundary between Layer 2 switching and Layer 3 routing. The aggregation layer provides Layer 3 services, routing LAN traffic between networks in and out of the campus. Because Layer 2 networks are terminated at the aggregation layer, it segments the network into smaller broadcast domains to reduce the size of the mac learning and IPv4/IPv6 tables for the local devices. 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, providing redundancy with VSX, high performance, and single hop connectivity between all switches in the aggregation block. The aggregation layer also becomes the ideal location for connecting other network services, such as the WAN aggregation, Internet DMZ, and server rooms for an organization.

Core Layer

In a large LAN environment, there are often multiple aggregation layer switches. When access layer switches are located in geographically dispersed buildings, costly fiber-optic runs between buildings are minimized 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 deploy a collapsed core which combines 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.

Nonstop Connectivity

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 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 all 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.

Campus Wired Summary

The ESP Campus wired LAN provides network access for employees, APs, and IoT devices. The campus LAN also becomes the logical choice 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.
- 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 connects 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.

Campus Wireless Architecture

Access Points are the underpinning of the ESP Campus wireless architecture. To provide maximum flexibility, there are two deployment modes an AP can operate in, bridge mode and tunnel mode. In addition to Wi-Fi 5 and Wi-Fi 6 connectivity, Aruba APs have Bluetooth and IEEE 802.15.4 (ZigBee) radios to provide wireless connectivity for IoT devices. While many vendors claim similar capabilities based on hardware capable of using IoT protocols, they haven't developed the deep integration at the software layer. Through tight IoT partnerships, Aruba has developed integrations with technologies such as digital door locks, electronic shelf labels, location tracking and much more using their close relationships to ensure successful deployments. By unifying the connectivity options, Aruba can orchestrate the use of airtime and reduce unneeded third party IoT Gateways. Also, to enable proprietary wireless protocols and allow expansion to an even wider variety of sensors, Aruba has enabled the USB port on APs for additional IoT connectivity.

Traditionally, enterprise class campus networks relied on controllers for data-plane centralization, wireless control-plane functionality, and the management of APs. With the Aruba ESP Architecture, Aruba moved the control-plane and management of the Gateways to Central for maximum scalability. This approach delivers consistent control-plane and management features between the two AP modes.

Organizations can optionally deploy a Gateway if they want to centralize the data-plane, apply advanced segmentation rules, or deploy UBT. These data-plane functionalities are applicable to both wired and wireless traffic. To ensure resiliency and the horsepower needed to handle the largest networks, horizontal-scale Gateway clustering was developed. With support for twelve node clusters and up to 100Gbps of stateful firewall throughput per node, the Aruba ESP Architecture can scale to meet the needs of any organization.

The wireless design provides reliable network access for employees, simple Internet access for visitors, and always-on 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 design 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 cloud-based control of wireless environment allows easy management and operation.
- Pay as you grow with gateway clustering for increasing network capacity and seamless failover.
- Live upgrades perform operating system updates, and in-service module updates support 24/7 operations.

Aruba wireless networks 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 environments. Each wireless 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.

AP Wireless

Like the wired models, the campus wireless models build on each other. The AP wireless model provides an easy solution when tunneled traffic is not needed, and the advanced Gateway features are not required. In this model, wireless traffic is bridged directly at the AP into the wired infrastructure. The access switch ports for the APs are trunked to connect the VLAN from each SSID into the switch. The AP handles the packet encryption, user authentication and policy enforcement functions, while features like RF management, key management, live upgrades, monitoring, and troubleshooting are managed in Central.

The switches handle Layer 2 adjacency to other clients so often customers will deploy this wireless option in conjunction with two-tier wired networks or switching networks with Layer 2 overlays. The following figure shows the AOS 10 AP wireless option.

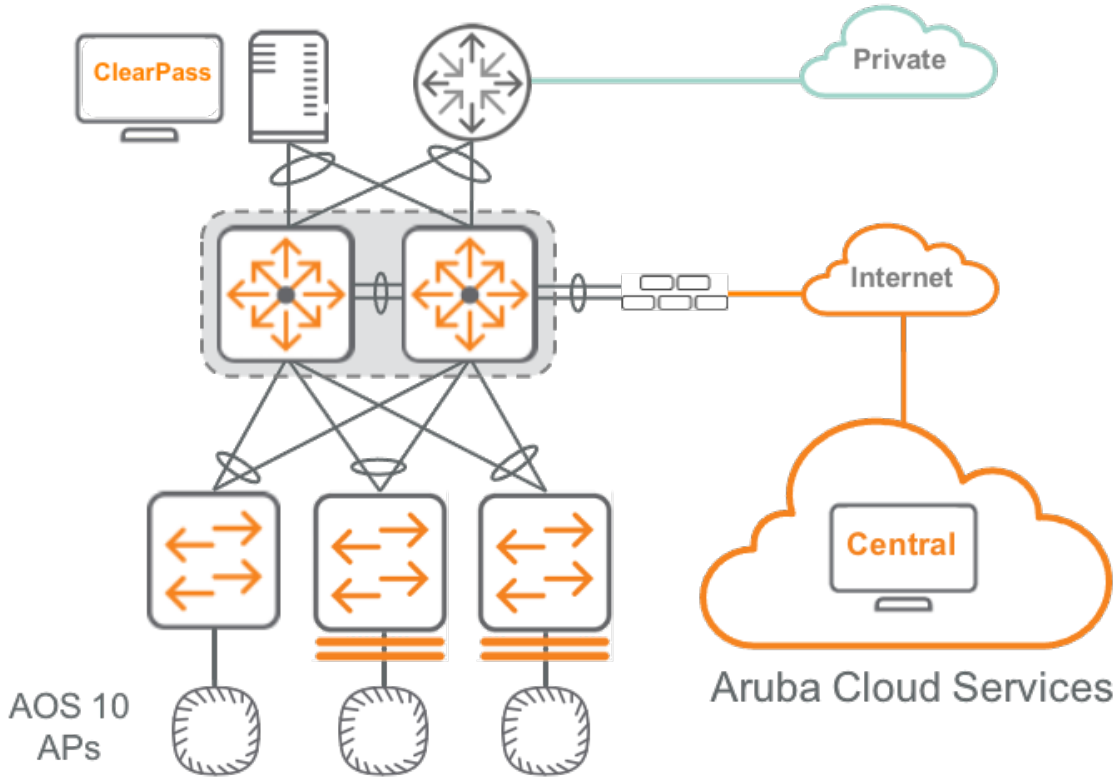

Figure 9: AP wireless

AP with Gateway Wireless

If an organization needs Gateway features, they can be added to the AP design using the same APs running the same software. The AP with Gateway option offers robust security features and maximum operational flexibility. Gateways can be deployed by themselves or clustered for additional redundancy and scalability. Clusters are automatically created by adding Gateways to a single group, and the tunnels between them are orchestrated in Central.

Like the AOS 10 AP option, this model includes automatic RF management to ensure the best Wi-Fi connections and granular visibility into applications, which helps prioritize business-critical data. This model also includes micro-segmentation, dynamic RADIUS proxy, and encryption over the LAN. Seamless roaming is supported across an entire Layer 3 campus. The following diagram shows the AOS 10 AP with Gateway wireless option.

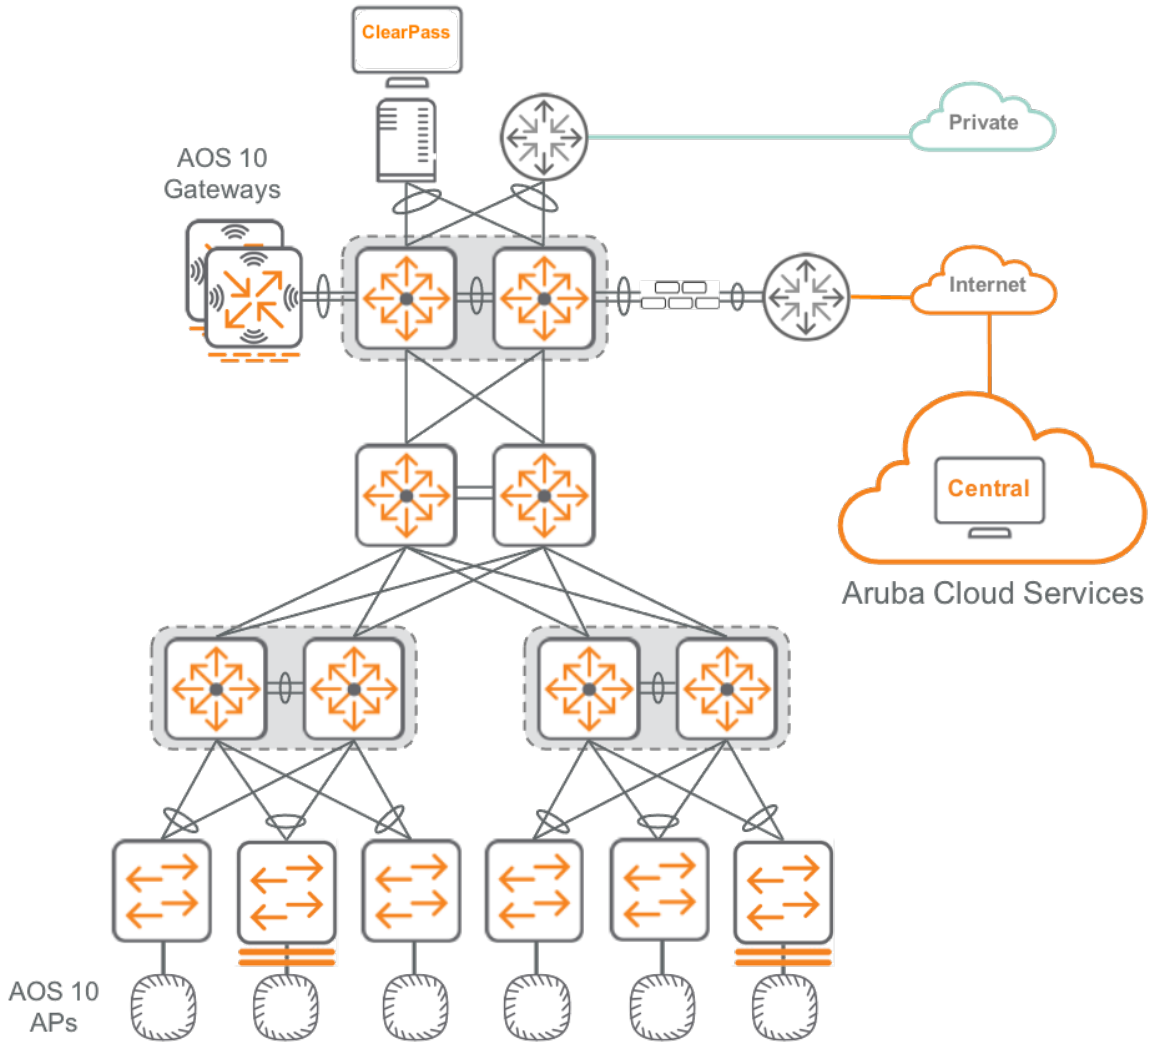

Figure 10: AP with Gateway wireless

Wi-Fi 6 Enhancements

Designed to address connectivity issues for high density deployments, the new Wi-Fi 6 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 Wi-Fi 5.

The most significant new feature of the Wi-Fi 6 standard is orthogonal frequency-division multiple access (OFDMA), which replaces orthogonal frequency-division multiplexing (OFDM). Other important new features include Basic Service Set (BSS) coloring and the ability to transmit up to 8 clients with Multi-User Multiple Input Multiple Output (MU-MIMO). Aruba recommends moving wireless clients to Wi-Fi 6 since both the client and APs have to be Wi-Fi 6 capable to take advantage of these new features.

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 the behavior and channels are divided into smaller sub-channels allowing the AP to 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 additional lower-speed clients. Sub-channel use is dynamic and adjusts automatically every transmission cycle, depending on client data needs.

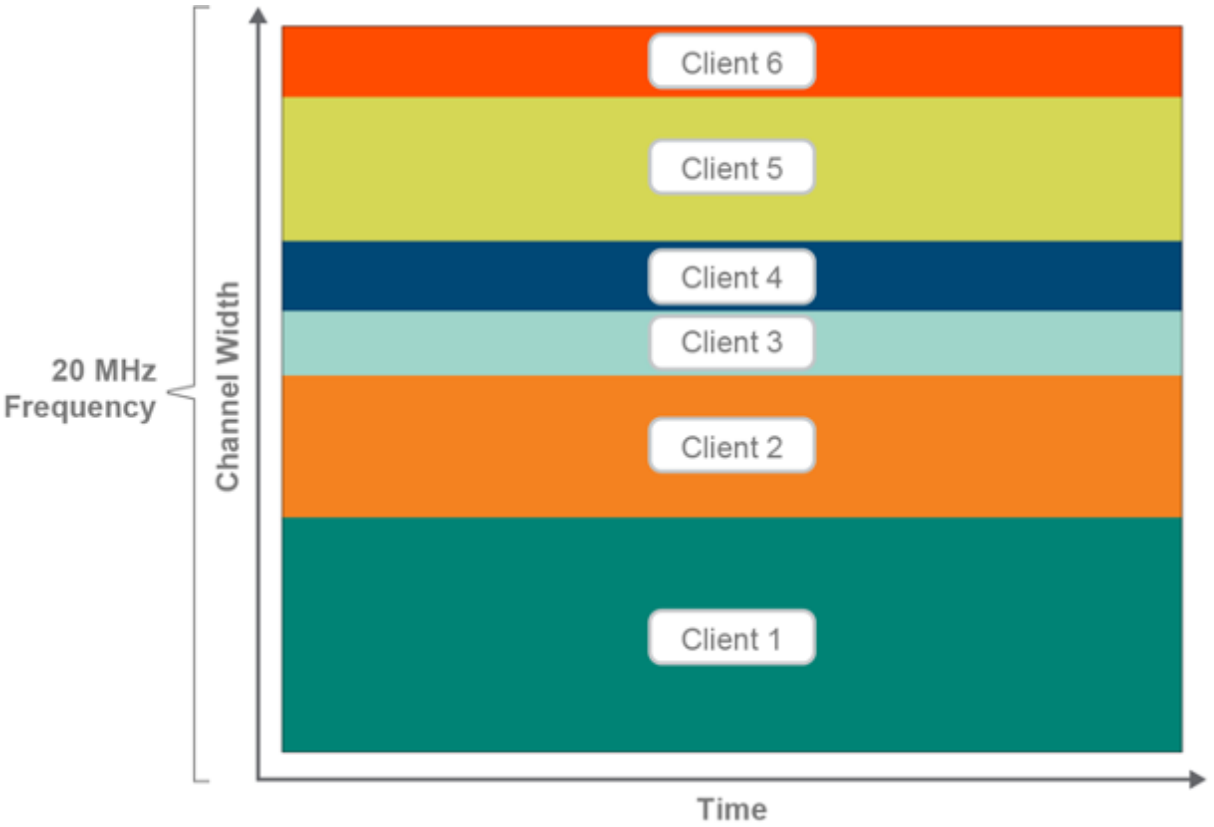

Figure 11: Multiple clients share the channel

This feature is enabled by default for Wi-Fi 6 clients and APs but only works when both sides are Wi-Fi 6 capable. Aruba recommends keeping it enabled as it leads to increased efficiency and reduced latency. Wider channels can support even more sub-channels which means 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.

Downlink MU-MIMO

The Wi-Fi 6 standard enhances MU-MIMO and supports up to eight clients simultaneously when using an eight Spatial Stream (SS) AP, like the Aruba 55X models. Even though the Wi-Fi 5 standard allowed for eight SS, most vendors implemented four or less, so this feature effectively doubles the number of devices which can be sent data from the AP at any given time. Having an increased number of spatial streams has the following benefits:

- Achieving higher data rates while communicating with a single client.
- Achieving higher aggregate performance in an MU-MIMO environment when simultaneously communicating with multiple clients.

That's the reason why AP-55x supports 8x8:8 MIMO, which equals 8 transmitting antennas, 8 receiving antennas, and 8 spatial streams.

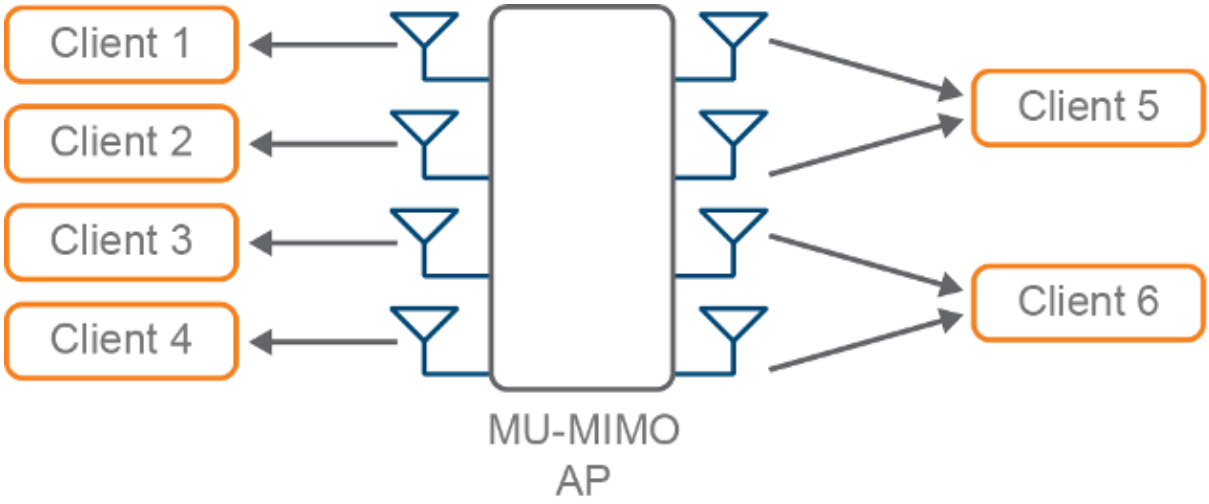

Figure 12: Single and dual stream clients

This feature is enabled by default and Aruba recommends keeping it enabled as it leads to increased capacity and higher speeds per user.

High Efficiency

All the Wi-Fi 6 specific features fall under the High-Efficiency profile. Enabling this parameter activates all the Wi-Fi 6 features on the radio. High Efficiency is enabled by default and Aruba recommends keeping it enabled.

Transmit Beamforming

Wi-Fi 6 employs an explicit beamforming procedure, similar to that of Wi-Fi 5. Under this procedure, the beamformer AP initiates a channel sounding procedure with a null data packet. The beamform client measures the channel and responds with a beamforming feedback frame containing a compressed feedback matrix. The beamformer uses this information to compute the antenna weights to focus the RF energy for each user. It is recommended to keep this feature enabled to achieve optimal performance benefits.

Target Wake Time

One of the main power saving features that Wi-Fi 6 offers is Target Wake Time (TWT), which is very useful for IoT devices. TWT uses negotiated policies based on expected traffic activity between Wi-Fi 6 clients and a Wi-Fi 6 AP to specify a scheduled wake time for each client. IoT clients that support Wi-Fi 6 could potentially sleep for hours/days and conserve battery life. This improves wake and sleep efficiency on smartphones and other mobile devices. Keep this feature enabled to allow clients to request a specific wake up time interval from the AP, thus allowing clients to sleep longer and save more power.

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 and channels with the same numbers in different colors do not overlap with each other, so four or five 80 MHz channels in an office environment would work. In the following example, the closest overlapping channels are the two sets of gray channels which are spaced far enough apart to not be a problem.

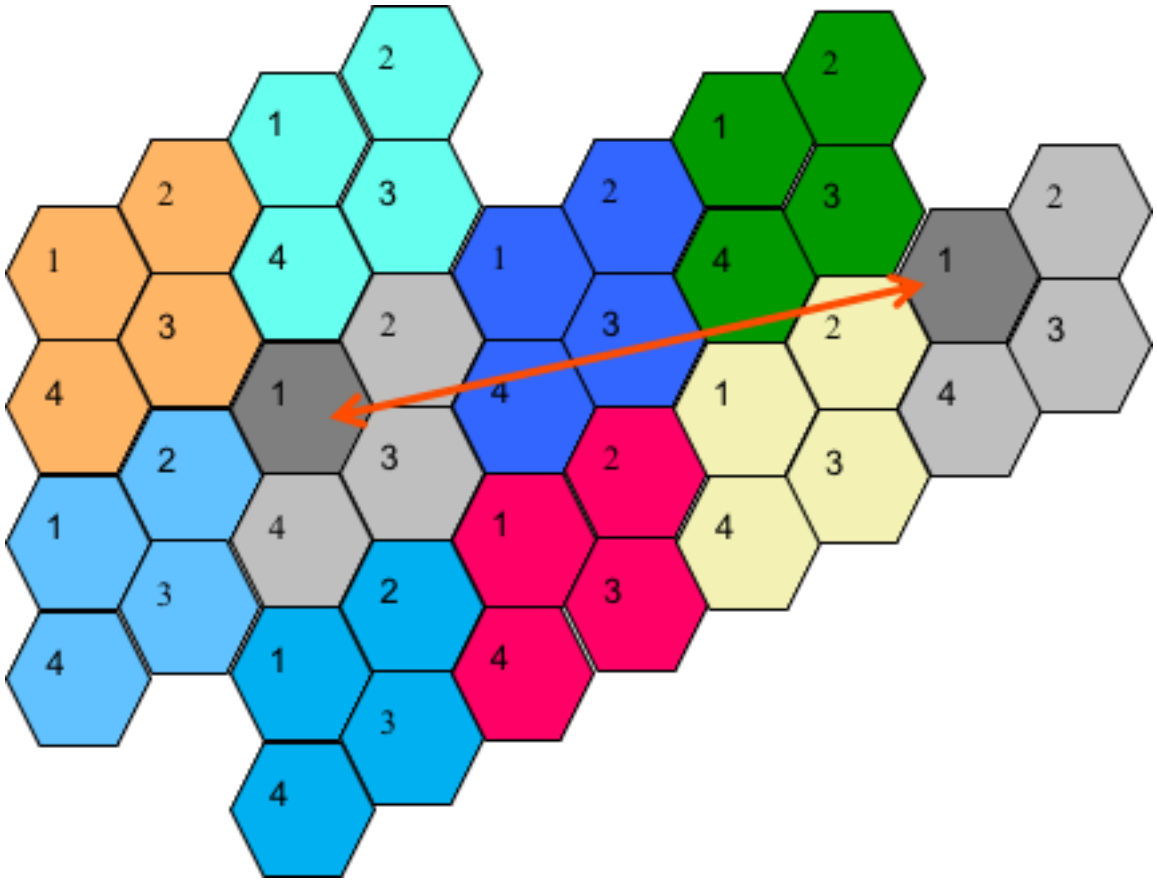

Figure 13: Same channel only blocked on color match

Wireless Security

Wireless security is a key aspect of the Wi-Fi 6 solution from Aruba. With all of the issues uncovered over the years related to WPA2, PSK Mode, and 802.1X/EAP, an effort was made to design suitable protocols that would help address each issue. Eventually, a family of protocols was defined that would become a new Wi-Fi Alliance certification program called WPA3. Aruba was instrumental in defining the protocols and bringing them to the Wi-Fi Alliance for full certification. Aruba recommends moving wireless clients to WPA3 as quickly as possible.

WPA3

Aruba Simultaneous Authentication of Equals (SAE) protocol was added in the late 2000s to the IEEE 802.11s mesh networking standard and 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 where each side proves it knows the password without exposing the password or any password-derived data. WPA3 can be deployed using WPA3-Personal (SAE) or WPA3-Enterprise.

NOTE:

The Protected Management Frame (PMF), otherwise known as 802.11w, is mandatory for organizations using WPA3 and OWE.

Please refer to the following URL for Wi-Fi Alliance's published list of client devices that are WPA3 certified: Wi-Fi Alliance WPA3 Certified Client Devices

WPA3-Personal

WPA3-Personal is a replacement for WPA2-PSK. As mentioned above, it uses password-based authentication based on Dragonfly key exchange (RFC 7664), which is resistant to active, passive, and dictionary attack. For backward compatibility, you must enable "Transition Mode" so that WPA3 capable client connects using WPA3-SAE, and legacy clients connect using WPA2-PSK.

NOTE:

ESP supports WPA3-Enterprise Basic as well as WPA3-Enterprise Suite-B deployment modes.

WPA3-Enterprise Basic

- This Operation Mode (opmode) is essentially the same as WPA2-Enterprise but with additional PMF settings.
- It is backward compatible with WPA2 and adds 802.11w support which is optional.
- Clients that are PMF capable (support 802.11w) and legacy clients (do not support 802.11w) can connect to same SSID. Hence, Transition Mode is not required for backward compatibility.
- This opmode is supported in, bridge, tunnel, mixed mode SSIDs.
- Recommended for enterprise customers even with non-CNSA strong encryption (GCM-256).

WPA3-Enterprise Suite-B

- Brings 192-bit security suite that is aligned with Commercial National Security Algorithm (CNSA) for enterprise network.
- In this opmode TLS cipher suites used in EAP-TLS are required, and PMF is mandatory.
- Recommended opmode for government agencies.
- This opmode is supported in, bridge, tunnel, mixed mode SSIDs.
- Client support for Suite B is not yet commercially available from Microsoft, but supported on MacOS

Enhanced Open

Aruba Opportunistic Wireless Encryption (OWE) provides unauthenticated data encryption to open Wi-Fi Networks. To the user, an OWE network looks just like an open network with no padlock symbol. However, the advantage is the data is 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 the client and the infrastructure. The key can be used to derive keys to encrypt all management and

data traffic sent and received by the client and AP. Key exchange happens on the initial AP, but then it is shared with the Key Management System (KMS) in Central and pushed to the neighboring APs by AirMatch and the KMS. The process is transparent to users and administrators, so additional device provisioning is not required for OWE.

Aruba recommends enabling OWE for visitor networks where encryption is needed, but authentication is not required like in coffee shops, bars, schools, public venues, and stadiums.

For backward compatibility, *Transition Mode* is required which allows an administrator to configure a single open SSID. The AP automatically creates two BSSIDs with separate beacons when OWE is enabled.

- **BSSID 1** - An open network for non-OWE stations which has an information element (IE) to indicate a BSSID 2 is available. Legacy clients connect to this BSSID; their traffic will not be encrypted.
- **BSSID 2** - A hidden OWE network with the Robust Secure Network Indicator Element (RSN-IE) Authentication Key Management (AKM) field indicating the use of suite 18 (the OWE suite) for authentication. In addition, an IE to indicate BSSID 1 is also available. OWE capable clients connect to the hidden SSID and will get Protected Management Frames (PMF) and encryption benefits.

Aruba supports configuring OWE SSID in bridge or tunnel mode. It is supported in the mixed mode SSID, but this deployment requires VLAN assignment rules based on RADIUS VSA attributes.

WPA3 and Enhanced Open represent a long overdue evolution for wireless security. The Internet and how it is used has changed considerably since WPA2 was released, and the problems associated with it have come to the forefront. WPA3 addresses the shortcomings of WPA2 and also addresses use cases not available with WPA2. An important part of WPA3 includes an increase in security while complexity remains the same. Increases in security are normally accompanied by large increases in complexity, which makes security harder to obtain and implement. The advantage of using WPA3 is there are no changes in workflows or usage, no new steps or caveats to remember. OWE looks just like an open network and allows the user to simply click and connect. WPA3-SAE looks just like WPA2-PSK, where a user enters their password and connects.

Multiple Pre-Shared Key

Multiple Pre-Shared Keys (MPSK) are used for IOT devices that are headless and unable to support network security functionality like 802.1X. MPSK takes advantage of new standards such as WPA3 and OWE to overcome the WPA2 pre-shared key problems by enabling device-specific and group-specific passphrases, which enhances security and deployment flexibility for headless IoT devices. Passphrases are administratively assigned to groups of devices based on common attributes like profiling data or uniquely assigned to each device registration with ClearPass Policy Manager.

MPSK provides the following benefits:

- Ensure that IoT devices are authenticated and legitimately connected to the network, without administrative involvement

- Several MPSKs are supported on a single SSID which improves the RF bandwidth utilization for all wireless users
- Establishes a one-to-one relationship between devices and a specific user to provide visibility, accountability and management
- Associates a device with a group of users, for example, a smart TV that is used by the marketing team

NOTE:

MPSK is not supported with SAE.

Visitor Wireless

Organizations often have a wide range of visitors that request network access while they are on-site. Visitors 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 visitors and their specific roles, an organization should deploy visitor access throughout the network and not only in lobby or conference room areas.

The flexibility of the Aruba ESP architecture allows the wireless network to provide visitor 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 visitor 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 visitor network:

- Provides Internet access to visitors through an open wireless Service Set Identifier (SSID), with web access control in the firewall.
- Supports the creation of temporary visitor authentication credentials that are managed by an authorized internal user.
- Keeps traffic on the visitor network separate from the internal network in order to prevent a visitor from accessing internal resources.
- In the future when more client devices support WPA3 and OWE, the visitor network will automatically encrypt traffic for the users.

Every AP can be provisioned with controlled, open access to wireless connectivity and the Internet. From the wireless AP, visitor traffic is securely tunneled back to the Gateway 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 visitor network VLAN to the firewall protecting the organization's private assets, as depicted in the following figure.

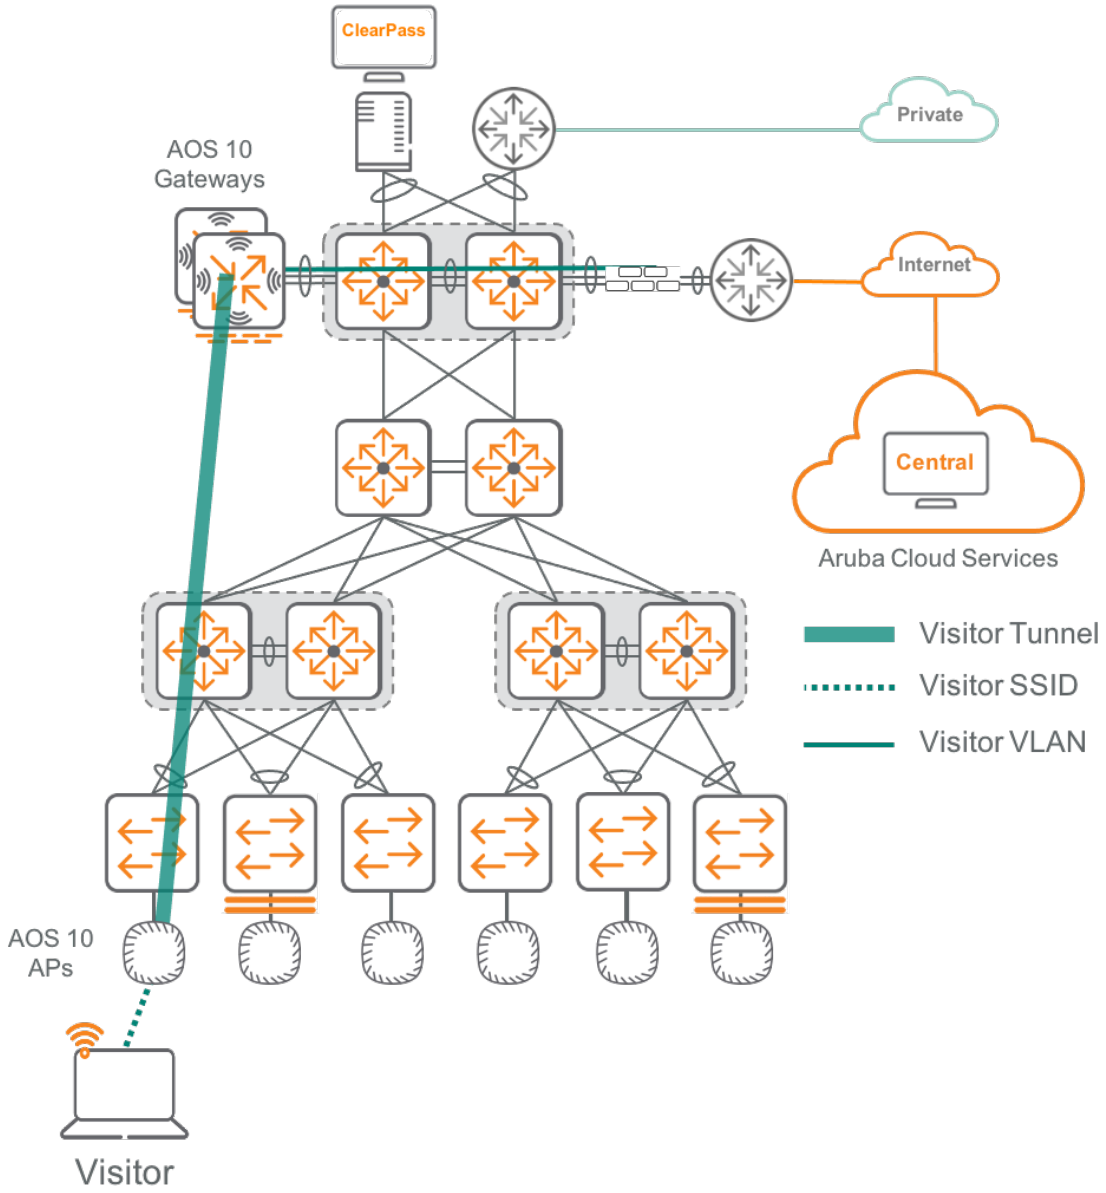

Figure 14: Visitor wireless network

To control connectivity, visitors are redirected to a captive portal and must present a username and password to connect to the visitor network. The captive portal can be on the Gateway or an external device. Because the visitor traffic must pass through the firewall, strict rules are applied to prevent visitor access to the internal corporate network. Lobby ambassadors or other administrative staff can assign temporary visitor 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.

It is common for the Gateway to act as a DHCP server and router for visitor clients especially when in a DMZ. As with all Gateway installations, as long as the metrics are below the recommended scaling capacity of the Gateway, it is acceptable to enable Layer 3 operations for use cases like Visitor or IOT. When the Gateway has routing enabled, firewall policies must be used to control traffic in and out of VLANs. When using the Gateway in a Layer 2 only mode, there is less of a concern as the external router or firewall will control inter-VLAN traffic. The DHCP service on the Gateway does not have redundancy, so if advanced DHCP services are required, an external DHCP server is recommended.

Campus Wireless Summary

The ESP campus WLAN provides network access for employees, visitors, 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, visitors and IoT devices.
- Plug and play deployment for wireless APs.
- Wi-Fi 6 enhancements to address connectivity issues for high density deployments and improve the performance of the entire network.
- Live upgrades to perform operating system updates with little to no service impact.

Live Upgrade

To ensure maximum uptime in the connectivity layer, Aruba has developed capabilities across both the wired and the wireless which orchestrate the process of software upgrades called Live Upgrade. When a Live Upgrade is triggered in AOS-CX, a pair of redundant switches go through multiple stages to minimize impact to network. First, the secondary switch installs the new software. Next, the secondary switch moves traffic to the primary by draining it's MC-LAG links and finally, the secondary switch reboots to initialize the new software. Once the secondary switch has upgraded and restored operations, the primary switch goes through the same process of installing software, moving traffic and upgrading. This method reduces lost traffic to nearly zero and can be triggered by one command in Central.

Wireless Live Upgrade accomplishes the same goal of nearly zero traffic loss but does so in a manner appropriate for wireless architectures. In this section the term Partition will be used to designate a group of APs that can be upgraded together with little to no downtime. The term Partition used here does not refer to the Partition on the Gateway or AP that will have code loaded on it. APs are Partitioned with data obtained from the AirMatch service. Traffic is moved so that all traffic from a specific AP Partition or Gateway has no active clients. This process happens through an integration between the Gateway clustering technology and ClientMatch on the APs. Once the traffic has completely drained from the devices, the AP or Gateway is upgraded. Once the process is completed the next Partition or

Gateway is upgraded and this process continues until all APs and Gateways are done. Live upgrade is available for a clustered pair of Gateways and APs that are deployed with overlapping coverage. The APs and Gateways can be upgrade separately and do not need the same code level as in previous versions.

ESP Campus Summary

Today's enterprise networking requirements are daunting. The exploding growth of IoT and mobility is proliferating the number of devices an organization must support. Legacy network architectures simply cannot scale to deliver the needed capabilities for configuring and provisioning with automation. Furthermore, IT operators demand a consolidated approach to managing connectivity across multiple domains and users want a consistent experience both inside and outside of the office walls to maintain a high level of productivity.

With Aruba's Campus Design, the most advanced network hardware and policy management is available to allow an organization to have a secure, agile, and automated network to support business needs. Aruba has structured ESP with a unified operating model that is IoT-ready and layered with flexible policy to utilize existing segmentation yet allows for an easy migration of a network to a modern cloud-services architecture. Cloud services provides elasticity and a unified data set for analytics. Aruba is standards based and works with a robust ecosystem of partners that ensure the latest innovations are available in a timely manner, no matter the size of an organization.

Campus Connectivity Layer

The *Connectivity Layer* is the foundation for the Aruba ESP architecture. It features the physical infrastructure components with extensive flexibility and high availability, along with the telemetry for analyzing the current state of the network. Enterprise networks need to support their existing endpoints, including many legacy systems, while transitioning to a modern architecture. This section will discuss the different design aspects of the Connectivity Layer and how they relate to the ESP Architecture.

OSPF 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. The two most common choices for an enterprise network are Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP). In this design, we recommend using OSPF for its simplicity and ease of configuration. 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.

In the three-tier and two-tier campus designs, the access switches have a default gateway in the management VLAN for operational access, and the VLANs are terminated at the aggregation layer switches. Aggregation layer switches are prevented from forming adjacencies with access switches, but they will advertise the available VLANs to the switches. The aggregation and core switches in the three-tier design use OSPF to dynamically exchange routes for the campus network. OSPF uses point to point links between aggregation and core devices for more efficient transfer of OSPF updates. The collapsed core and WAN gateways in the two-tier design use OSPF to dynamically exchange routes for the campus network.

Enabling authentication for OSPF devices is only required for highly secure networks, for all other network deployments authentication is optional. The Internet is accessed using a static default route originating from the DMZ firewall. The diagram below outlines the key points for running OSPF in a three-tier wired design.

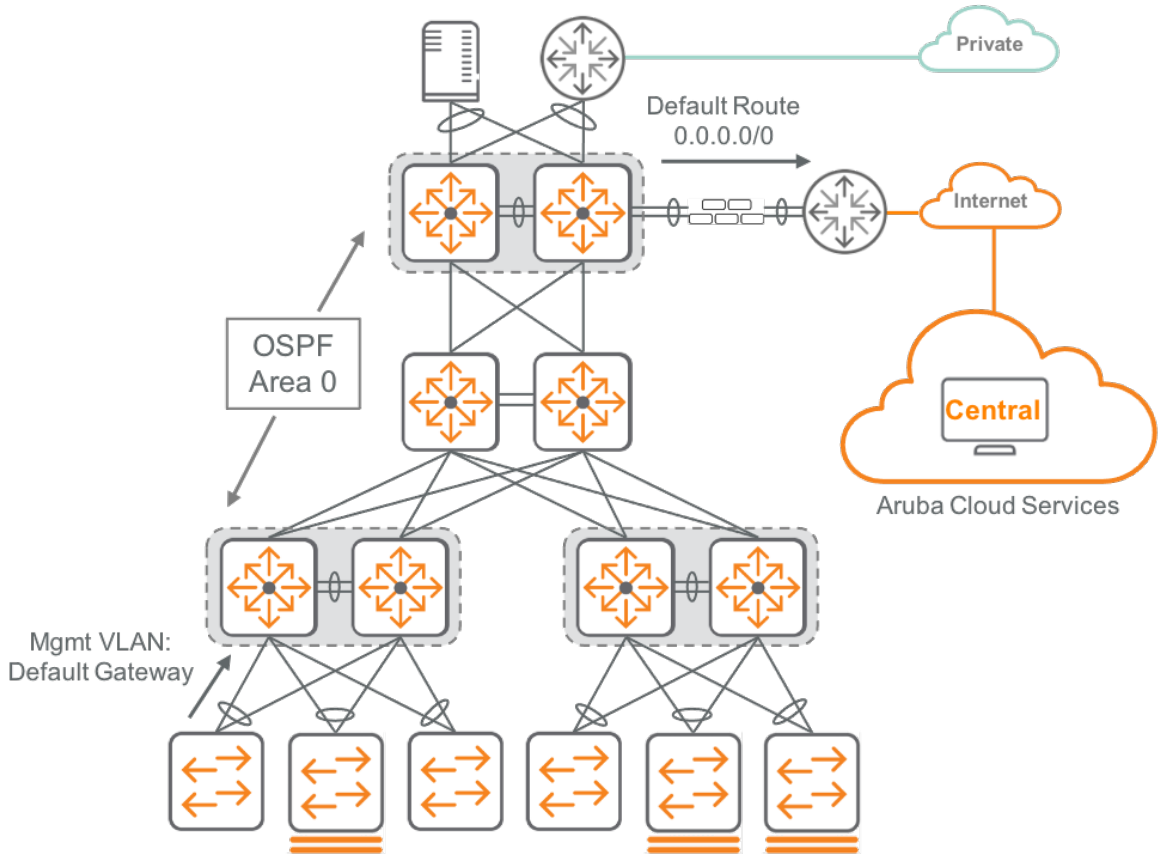

Figure 15: OSPF in three-tier wired

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 of pre-recorded content.

This design uses protocol independent multicast-sparse mode (PIM-SM) to route multicast traffic on the network. The mechanisms to route multicast traffic are rendezvous point (RP), bootstrap router (BSR), Multicast Source Discovery Protocol (MSDP) and Internet Group Management Protocol (IGMP). PIM-SM should be configured on all routed links to enable multicast traffic on the network. In this design, the OSPF routing table is used for reverse path forwarding to direct the multicast traffic.

The RP is the root of the multicast tree when using sparse mode. Multiple RPs can be configured for redundancy, although normally, only one RP is active at a time for each multicast group. Multiple RPs can be active if MSDP is enabled because it allows a multicast domain to share source tree tables between RPs. MSDP allows switches to have Inter and Intra Domain active-active redundancy using an Anycast IP address as the RP. Anycast is a networking technique that allows for multiple devices to share the same IP address. Based on the location of the user request, the switches send the traffic to the closest device in the network which reduces latency and increases redundancy.

In a Campus, MSDP is needed for intra domain redundancy and should be enabled on the core switches in either the two-tier or three-tier topologies. The RP candidate announcement, in combination with MSDP, advertises the Anycast IP address to neighboring devices. Neighboring devices will not know what devices want to be the RP unless BSR is enabled. The BSR is elected from a list of candidate-BSRs configured on the network. There can only be a single active BSR, and it advertises RP information to all PIM-enabled routers, freeing the administrator from having to statically configure the RP address on each router in the network. BSR, RP and MSDP should be enabled on the core switches to identify the active RP and notify neighboring devices.

When a client wants to join a multicast group, it sends an IGMP join message to the local multicast router which is also known as the designated router (DR). The DR forwards the join message towards the RP and all routers in the path do the same until the IGMP join message reaches the RP. Multicast traffic is forwarded back down the shared tree to the client. Periodic IGMP 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 an IGMP 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. IGMP should be enabled on aggregation switches and collapsed core switches.

NOTE:

IGMP timers must match across all switches on the network, including switches from other vendors.

Dynamic Multicast Optimization

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. Aruba recommends enabling Dynamic Multicast Optimization (DMO) to allow the AP to convert the multicast traffic to unicast for each client device. DMO is a technology from Aruba that converts multicast frames over to the air to unicast. This provides the same benefits as ARP optimization to decrease the channel duty cycle and guarantee frame delivery. Unicast frames are acknowledged by the client and can be retransmitted if a frame is lost over the air. Unicast frames are also transmitted at the highest possible data rate supported by the client vs the lowest basic rate.

Multicast and broadcast frames are natively transmitted at the lowest basic rate in order to have a higher chance of successful delivery to all associated clients. With DMO enabled, the multicast packet is converted in the AP/Gateway datapath. This operation can be CPU intensive depending on the multicast packet size and number if multicast streams. It is not recommended to have multiple multicast sources with the same data crossing the datapath with DMO enabled. If multicast is required, it is recommended to use the largest possible Layer 2 network so multiple multicast streams are not converted at the same time. If large Layer 2 networks cannot be created, it is recommended to use additional Gateways just for the conversion. Enabling DMO in a properly sized installation does not have a negative impact on the Gateway or AP performance.

The following figure shows IP multicast BSR, RP, MSDP, IGMP Snooping and DMO placement in the two-tier and three-tier wired designs.

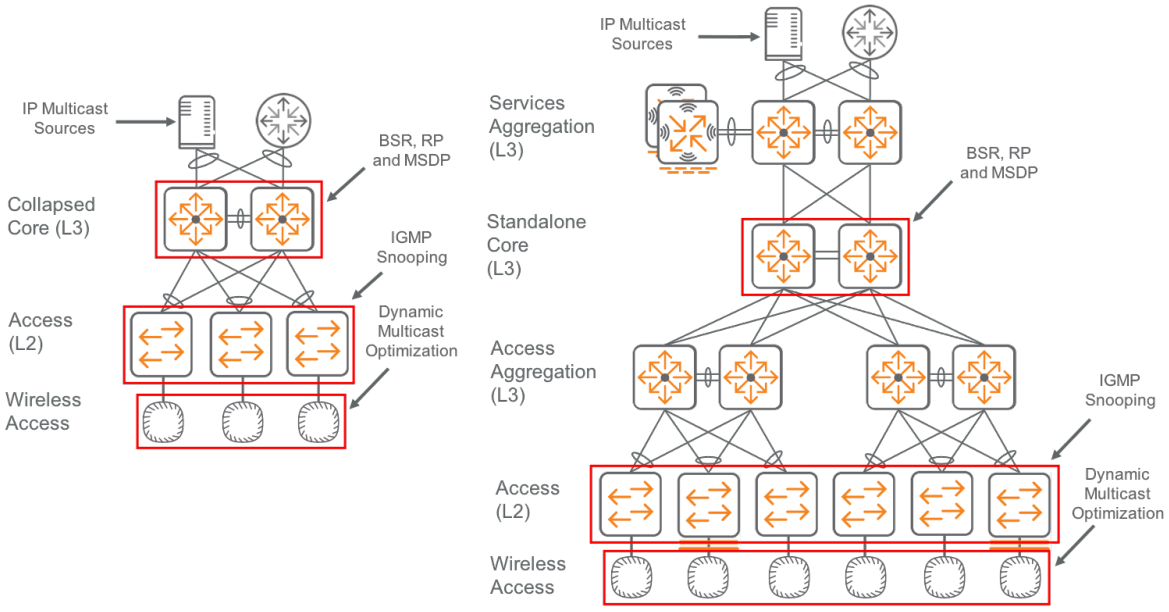

Figure 16: IP Multicast BSR, RP, MSDP, IGMP Snooping and DMO placement

Broadcast to Unicast Conversion

Aruba WLANs also can convert broadcast packets into unicast frames over the air. The advantage is because broadcast frames over the air are required to be transmitted over the lowest possible data rate configured called the “basic rate” in an attempt to guarantee delivery of every frame. Since broadcasts have no acknowledgement of delivery like unicast does, there is not an option to retransmit a lost broadcast frame. When the frame over the air is converted to unicast, the AP can send the frame at a much higher data rate and get confirmation of delivery. Retransmitting lost unicast frames is possible because each frame is acknowledged.

Converting broadcast to unicast has two large benefits, first as described above, guaranteed delivery of the frame and second, it will greatly decrease channel duty cycle and deliver the frames at the highest possible data rate on a per client basis. When broadcast frames are delivered at the basic rate for every client, it is equivalent to a single lane road with someone driving under the speed limit. When there isn’t a lot of traffic it isn’t terrible but when there is a lot of traffic the AP duty cycle suffers.

ARP Optimization

There are two modes available for ARP optimizations, “ARP Only” and “All”. With “ARP Only” the AP will only optimize ARP requests and the rest of the broadcast and multicast traffic will be forwarded as usual. With “All” enabled, every multicast and broadcast will be dropped. In order for multicast to be converted to unicast please see Dynamic Multicast Optimization section above.

Quality of Service

Quality of service (QoS) refers to the ability of a network to provide higher levels of service using traffic prioritization and control mechanisms. Applying the proper QoS policy is important for real-time traffic, such as Skype or Zoom, and business-critical applications, like Oracle or Salesforce. To accurately configure QoS on a network, there are several aspects to consider, such as bit rate, throughput, path availability, delay, jitter, and loss. The last three—delay, jitter and loss—are easily improved by using an appropriate queueing algorithm which allows the administrator to schedule and deliver applications with higher requirements before applications with lower requirements. The areas of the network that require queueing are the places with constrained bandwidth, like the wireless or WAN sections. Wired LAN uplinks are designed to carry the appropriate amount of traffic for the expected bandwidth requirements and since QoS queueing does not take effect until there is active congestion, queueing is not typically needed on LAN switches.

The easiest strategy to deploy QoS is to identify the applications running in the network that are critical and give them a higher level of service using the QoS scheduling techniques described in this guide. The remaining applications stays in the best-effort queue to minimize the upfront configuration and to lower the day-to-day operational effort of troubleshooting a complex QoS policy. If additional applications become critical in the future, they are identified and added to the existing list of business-critical applications. This can be repeated as needed without requiring a comprehensive policy for all applications on the network. This strategy is normally used by organizations who do not have a corporate-wide QoS policy or are troubleshooting application performance problems in specific areas of the network.

An example of this type of strategy prioritizes real-time applications, along with a few other key applications that require fast response times because users are waiting on remote servers. Identifying business critical applications and giving them special treatment on the network allows employees to remain productive doing the tasks that matter the most to a business. Real-time applications are placed into a strict priority queue and business critical applications are serviced by one or more premium queues that provide a higher level of service during times of congestion. The administrator must be careful to limit the amount of traffic placed into strict priority queues to prevent over saturation of the interface buffers. If all traffic is marked with strict priority, the QoS strategy becomes nothing more than a large first-in first-out queue which defeats the purpose of creating a QoS strategy. After critical traffic is identified and placed into the appropriate queues, the rest of the traffic is placed into a default queue with a lower level of service than the applications used for running the business. If the higher priority applications are not using the bandwidth assigned, the default queue uses all the available bandwidth.

Traffic Classification and Marking

Aruba recommends using the access switch as a QoS policy enforcement point for traffic over the LAN. This means selected applications are identified by IP address and port number at the ingress of the access switch and marked for special treatment. Traffic is optionally queued on the uplinks, but this is not a requirement with a properly designed campus LAN environment. Any applications that are not identified are re-marked to value of zero, giving them a best-effort level of service. The diagram below shows where traffic is classified, marked and queued as it passes through the switch.

Figure 17: Classification, marking and queueing

In a typical enterprise network, applications with similar characteristics are categorized based on how they operate over the network. These application categories are sent into different queues according to the types of applications. For example, if broadcast video or multimedia streaming applications are not used for business purposes, there is no need to account for them in a QoS policy. However, if Skype and Zoom are used for making business-related calls, the traffic must be identified and given a higher priority. Certain traffic that is not important to the business, like YouTube, gaming and general web browsing, should be identified and placed into the scavenger class, which allows it to be dropped first and with the highest frequency during times of congestion.

A comprehensive QoS policy requires business relevant and scavenger class applications to be categorized on the network. Layer 3 and Layer 4 classifications group the applications together into categories to help identify the ones with similar characteristics. After sorting the applications that are important from the ones that are not, they are combined into groups for queuing. Scheduling algorithms rely on classification markings to identify applications as they pass through a network device. Aruba recommends marking applications at Layer 3, rather than Layer 2 so the markings are carried throughout the life of a packet. The goal of the QoS policy is to allow critical applications to share the available bandwidth with minimal system impact and engineering effort.

DiffServ Code Point (DSCP)

Layer 3 marking uses the IP type of service (ToS) byte with either the IP Precedence three most-significant bit values from 0 to 7 or the DSCP six most-significant bit values from 0 to 63. The DSCP values are more common because they provide a higher level of QoS granularity, but they are also backward compatible to IP precedence because of their left-most placement in the ToS byte. Layer

3 markings are added in the standards-based IP header, so they remain with the packet as it travels across the network. When an additional IP header is added to a packet, like in the case of traffic in an IPsec tunnel, the inner header DSCP marking must be copied to the outer header to allow the network equipment along the path to use the values.

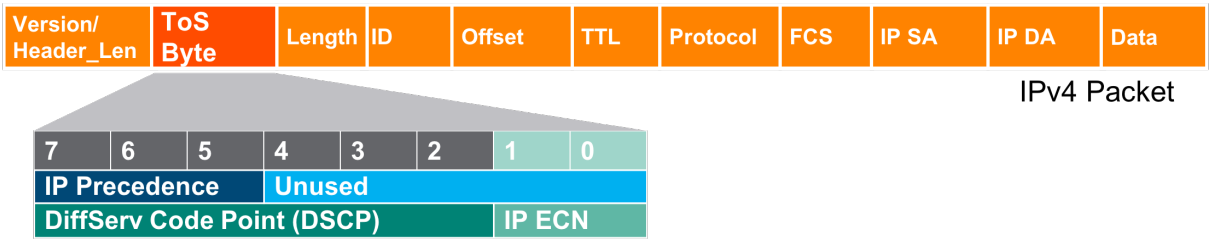

Figure 18: DSCP marking

There are several RFCs associated with the DSCP values as they pertain to the per-hop behaviors (PHBs) of traffic as it passes through the various network devices along its path. The diagram below shows the relationship between PHB and DSCP, along with their associated RFCs.

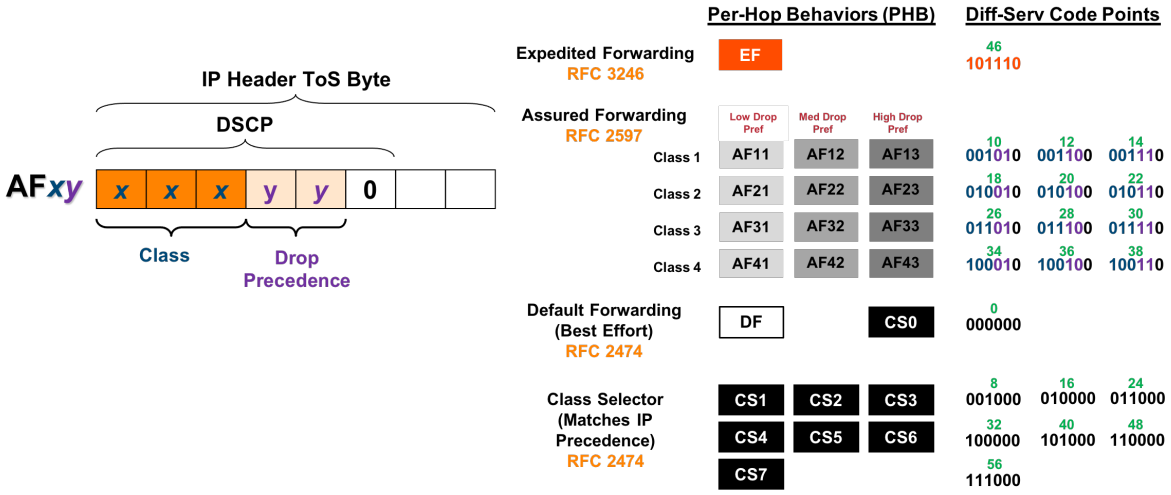

Figure 19: DSCP relationship with per-hop behaviors

Voice traffic is marked with the highest priority using an Expedited Forwarding (EF) class. Multimedia applications, broadcast and video conferencing are placed into an assured forwarding (AF31) class to give them a percentage of the available bandwidth as they cross the network. Signaling, network management, transactional and bulk applications are given an assured forwarding (AF21) class. Finally, default and scavenger applications are placed into the Default (DF) class to give them a reduced amount of bandwidth but not completely starve them during times of interface congestion. The figure below shows an example of six application types mapped to a 4-class LAN queueing model.

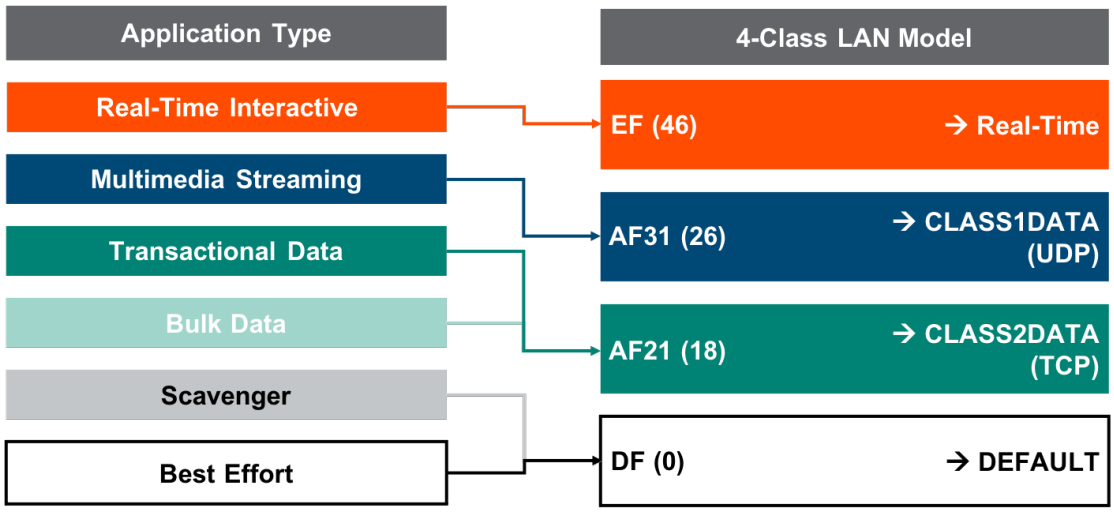

Figure 20: Six application types in a 4-class LAN queueing model

The "Best effort" entry at the end of the QoS policy marks all application flows that are not recognized by the Layer 3 / Layer 4 classification into the best effort queue. This prevents end users who mark their own packets from getting higher priority access across the network.

Traffic Prioritization and Queuing

The Aruba switch supports up to eight queues, but in this example only four queues are used. The real-time interactive applications are combined into one strict priority queue, while multimedia and transactional applications are placed into deficit round robin (DRR) queues, and the last queue is used for scavenger and default traffic. DRR is a packet-based scheduling algorithm that groups applications into classes and shares the available capacity between them according to a percentage of the bandwidth which is defined by the network administrator. Each DRR queue is given its fair share of the bandwidth during times of congestion, but all of them can use as much of the bandwidth as needed when there is no congestion.

The outbound interface requires the DSCP values shown in the second column in order to queue applications through the queues. The weighted values used in the DRR LAN scheduler column are added together and each DRR queue is given a share of the total. The values will need to be adjusted according to the volume of applications in each category on the network. The adjustment process is often done with trial and error as the QoS policy is used to affect the applications in the environment. The queues are sequentially assigned in a top-down order as shown in the 4-class example below.

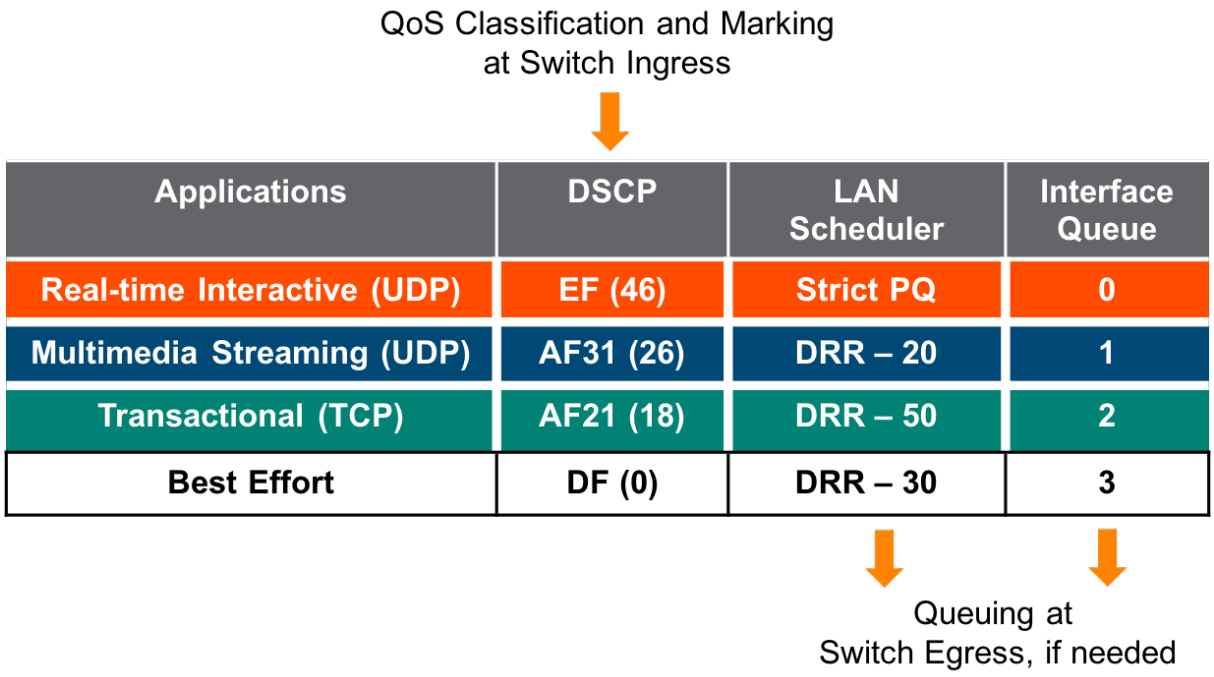

Figure 21: QoS summary for Aruba switch

Wi-Fi Multimedia

Wi-Fi Multimedia (WMM) - is a certification program that was created by the Wi-Fi Alliance that covers QoS over Wi-Fi networks. WMM prioritizes traffic into one of four queues, and traffic receives different treatment in the network based on its traffic class. The different treatment includes a shortened wait time between packets and tagging of packets using DSCP and IEEE 802.1p markings. Aruba allows users to define which traffic fits in to each queue and the DSCP and 802.1p values can be adjusted appropriately to match the network.

To take advantage of WWM functionality in a Wi-Fi network, three requirements have to be met:

1. The access point is Wi-Fi Certified™ for WMM and has WMM enabled.
2. The client device must be Wi-Fi Certified™ for WMM
3. The source application supports WMM.

NOTE:

WMM is supported in all of Aruba's Wi-Fi products.

QoS is not only set for a VLAN or port but can be set dynamically per application using Policy Enforcement Firewall. Most networks, including wireless LANs, operate far below capacity most of the time. This means there is very little congestion and traffic experiences good performance. QoS provides predictable behavior for those occasions and points in the network when congestion is experienced. During overload conditions, QoS mechanisms grant certain traffic high priority, while making fewer resources available to lower-priority clients. For instance, increasing the number of voice users on the network may entail delaying or dropping data traffic.

The Wi-Fi network is shared across multiple clients and the medium is bandwidth limited. The wireless spectrum occupied by an RF channel is shared by an access point, its associated clients, and by all other access points and clients in the vicinity that are using the same channel. Prior to Wi-Fi 6 and BSS Coloring, only one client or AP could transmit at any given time on any channel.

Wi-Fi uses carrier-sense, multiple-access with collision avoidance (CSMA/CA), much like the shared Ethernet networks did in the early days. Before transmitting a frame, CSMA/CA requires each device to monitor the wireless channel for other Wi-Fi transmissions. If a transmission is in progress, the device sets a back-off timer to a random interval and tries again when the timer expires. Once the channel is clear, the device waits a short interval – the arbitration inter-frame space – before starting its transmission. Since all devices follow the same set of rules, CSMA/CA ensures "fair" access to the wireless channel for all Wi-Fi devices. The Wi-Fi standard defines a distributed system in which there is no central coordination or scheduling of clients or APs.

The WMM protocol adjusts two CSMA/CA parameters, the random back-off timer and the arbitration inter-frame space, according to the QoS priority of the frame to be transmitted. High-priority frames are assigned shorter random back-off timers and arbitration inter-frame spaces, while low-priority frames must wait longer. WMM thereby gives high priority frames a much higher probability of being transmitted sooner. A station with low-priority traffic, on seeing another station transmit, must set the back-off timer to a random number within a broad range, say 15 to 1024. A station with high-priority traffic will select a random number from a smaller range, say 7 to 31. Statistically this ensures that the high priority frame will be transmitted with a shorter delay and has a lower probability of being dropped.

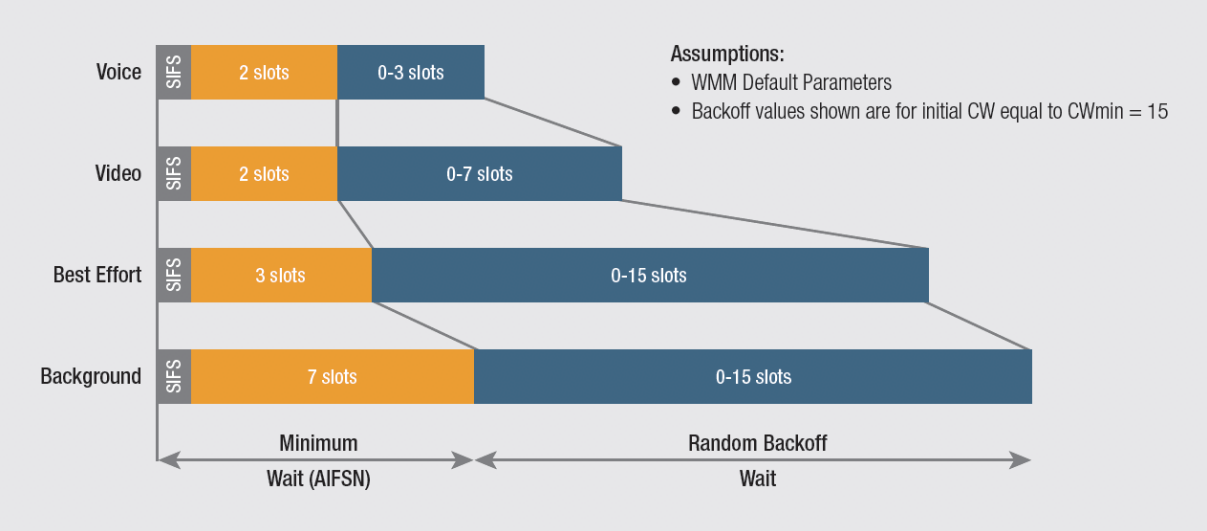

Figure 22: **Back-off and Arbitration inter-frame timers for WMM**

When a high-priority frame is served by a Wi-Fi network interface, the device is allowed to use a shorter arbitration inter-frame space than other devices using the same channel. This means that when the wireless channel goes quiet, devices with high-priority frames wait a shorter inter-frame space relative to other devices with lower priority traffic. This mechanism thereby assures more rapid transmission of high priority traffic.

The random back-off timer and arbitration inter-frame space mechanism address conditions during which multiple devices have traffic to transmit at the same time, and the total offered traffic is high relative to the capacity of the channel. However, these mechanisms don't address how a particular client or AP ensures QoS within its own interface during a temporary traffic peak. That capability is handled by an internal priority queuing mechanism. As packets are sent to the MAC layer of the Wi-Fi interface, they are internally lined up in their respective priority queues which are serviced in strict priority order. If the device generates more traffic than it can transmit onto the wireless channel, the higher priority traffic will override other packets within the interface.

WMM defines four priority levels in ascending priority: background, best effort, video, and voice. The default values for random back-off timers and arbitration inter-frame spaces are defined in the 802.11 standard, as is the queuing structure in the Wi-Fi interface. Since QoS must be maintained end-to-end, it is important that WMM priority levels be mapped to the QoS priorities in use on the LAN. The table below shows how DSCP priorities are translated into the four WMM priority levels.

Table 1: WMM to DSCP mapping

WMM Access Category	Description	DSCP
Voice Priority	Real-time interactive	46
Video Priority	Multimedia streaming	26
Best Effort Priority	Transactional	18
Background Priority	Best effort	0

Spanning Tree

High availability is a primary goal of any enterprise to conduct business on an ongoing basis. One method to ensure high availability, is to provide Layer 2 redundancy with multiple links. Without this redundancy, the failure of a link will result in a network outage. However, adding multiple links also introduces the potential for Layer 2 loops in the network. The spanning tree protocol (STP) can prevent loops in the network, regardless of the network topology. This section will dive into what devices should be root of the spanning tree topology, and what version of STP to use. This section will also cover supplemental features that should be enabled.

With many different versions of spanning tree and different network devices using different defaults, it is important to standardize on a common version of spanning tree running in order to have a predictable spanning tree topology. The recommended version of spanning tree for Aruba Gateways and switches is Rapid Per VLAN Spanning Tree (Rapid PVST+).

Spanning Tree and Root Bridge Selection

Spanning tree should be enabled on all devices as a heavy-handed loop prevention mechanism. This should be done regardless of network topology to prevent accidental loops. Gateways and access switches should have high bridge IDs to prevent them from becoming the root bridge of the network. Any Layer 3 device can be left at the default priority, as it is unlikely Layer 2 VLANs will be stretched across these devices so there is not a need to configure STP on them. The root bridge needs to be the device or pair of devices that are central to the network and aggregate VLANs for downstream devices. In the campus topologies discussed in this guide, the root bridge candidates are the collapsed core, access aggregation and services aggregation devices.

In the three-tier wired design, the root bridges are the access aggregation switches and service aggregation switches. As mentioned in the campus design overview, VSX and MC-LAG are used to allow dual-connections between the access and aggregation layers without the need for STP on the individual links. Even though there are multiple root bridges, they will not interfere with each other because spanning tree does not extend over the Layer 3 boundary between the devices. In this example, the access switches are Layer 2 and will need to be set with a high bridge ID. The core devices are Layer 3 switches and do not need to be set to a specific spanning tree value.

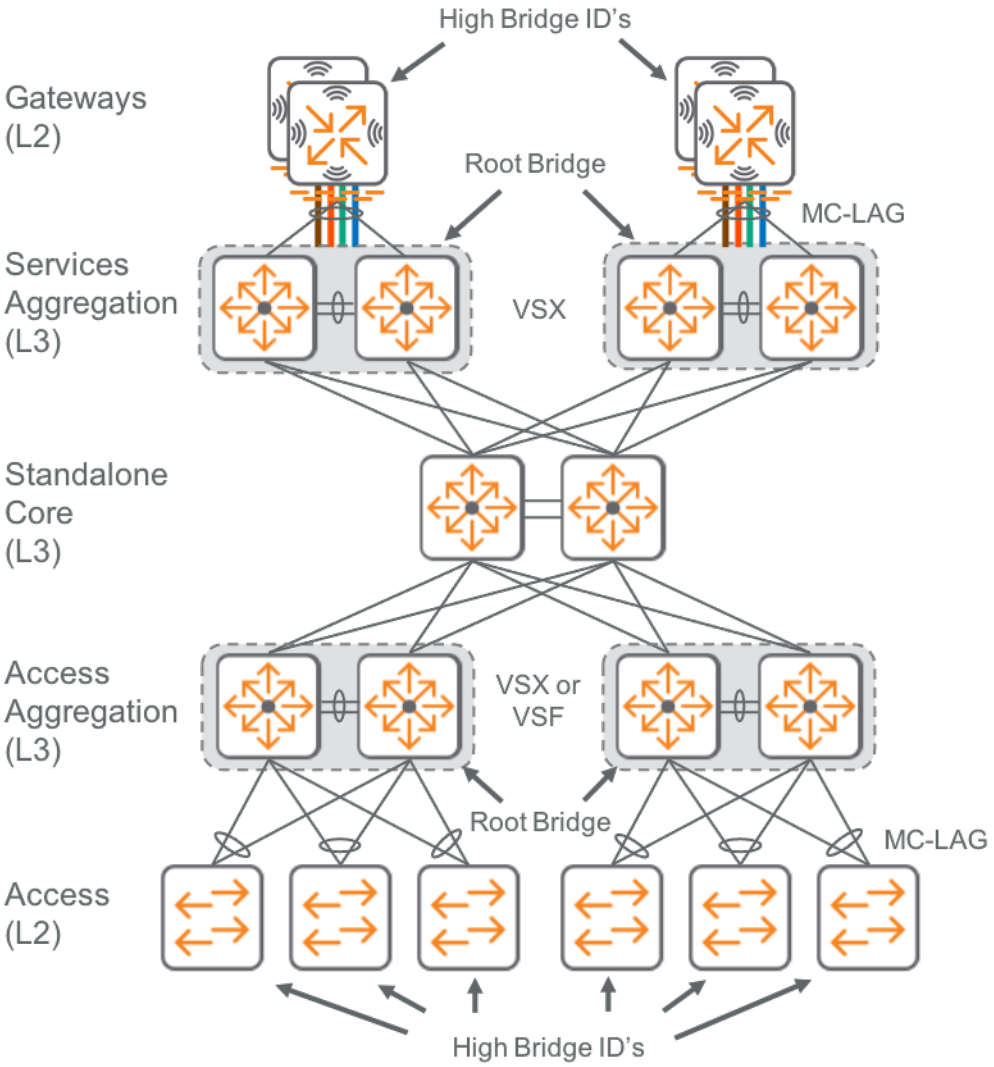

Figure 23: **Spanning tree root bridge placement**

Spanning Tree Supplemental Features

Spanning tree has several supplementary features that help keep the network stable. In this section, there is a brief overview of each feature along with a description of why they should be enabled.

Root Guard - Stops devices from sending a superior BPDU on interfaces that shouldn't send a superior or lower priority BPDU on an interface. This should be enabled on the aggregation or collapsed core downlinks to prevent access switches from becoming root of the network. This should not be enabled on the links connecting the aggregation switches to the core switch.

Admin Edge - Allows a port to be automatically enabled without going through the listening and learning phases of spanning tree on the switch. This is should only be used on single device ports or with a PC, daisy chained to a phone. This needs to be used with caution, because spanning tree will not run on these ports so if there is a loop on the network it will not be detected by spanning tree. This feature should only be used for client facing ports on access switches.

BPDU Guard - Automatically blocks a port if it sees a BPDU on them. This feature typically should be enabled on admin-defined client facing ports on access switches. This is to prevent any BPDU's from being received on ports that have been configured as a client facing port. BPDU guard ensures that BPDU's do not get received on access ports preventing loops and spoofed BPDU packets.

Loop Protect - Allows the switch to automatically block a loop if detected and unblock a loop automatically if the loop has disappeared. This feature should be turned on for all access port interfaces to prevent accidental loops from access port to access port. Loop Protect should not be enabled on the uplink interfaces of access switches or in the core, aggregation or collapsed core layers of the network.

BPDU Filter - Ignores BPDU's that are sent to an interface and will not send any of its own BPDU's to interfaces. The main use case for this feature is in a multi tenancy environment when the servicing network does not want to participate in the customers spanning tree topology. A BPDU filter enabled interface will still allow other switches to participate in their own spanning tree topology. It is not recommended to use BPDU filter unless there is a reason the network infrastructure does not need to participate in the spanning tree topology.

Fault monitor - This feature can be used to automatically detect excessive traffic and link errors. Fault monitor can be used to log events, send SNMP traps or temporarily disable a port. It is recommended to enable fault monitor in notification mode for all recognized faults. This feature should be enabled on all interfaces for data continuity, but do not enable the disable feature with fault monitor because spanning tree and loop protect is used to stop loops.

Radio Frequency Design

A site survey is an important tool for understanding the radio frequency (RF) behavior at a site and, more importantly, where and how much interference might be encountered with the intended coverage zones. A survey also helps to determine the type of wireless network equipment, 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 the network requires. Since RF interacts with the physical world around it, and because all office environments are unique, each wireless network has slightly different characteristics.

To provide ubiquitous multimedia coverage in a multi-floor/multi-building campus with uninterrupted service, the correct RF elements are required to ensure a successful implementation. Planning tools have evolved with the radio technologies and applications in use today, but a familiarity with the RF design elements and mobile applications are required to produce a good plan. Completing a site survey yields good information that can be used again and again as the wireless network grows and continues to evolve.

AirMatch

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 AirMatch feature provides automated RF optimization by dynamically adapting to the ever-changing RF environment at the network facility.

AirMatch enables the following key functions:

- Compute channel, channel width, and transmit power for APs
- Deploy channel and power plan based on configured deployment times
- Provide neighbor APs list to the Key Management service
- Provide AP data to the Live Upgrade service

In the ESP solution, the AirMatch service is moved to Central, which is capable of computing and deploying RF allocation to APs across the entire enterprise network. The AirMatch service receives telemetry data from APs for radio measurements, channel range, transmit power range, operational conditions, and local RF events like radar detection or high noise.

To determine the plan improvement, the average radio conflict metric is created. For each radio of an AP, the overlapped channels with its neighbors are calculated, and path loss weight is used to come up with the average weight value. After AirMatch comes up with a new plan, its conflict value is compared with the current operating network, and an improvement percentage is calculated. If the improvement percentage is higher than or equal to the configured quality threshold which is 8% by default, then a new plan will be deployed at the configured “Automatically Deploy Optimization” timer. The AP can still make the local channel changes in the case of poor service to a client. These localized channel changes are done without disturbing the entire channel plan. This information is relayed to the AirMatch service so the AirMatch engine can factor in the changes for future channel plans.

It is recommended to configure AirMatch wireless coverage tuning value to Balanced. When making changes to AirMatch, remember channel change events are disruptive, so it should be done only when absolutely required. However, AirMatch can only optimize the RF environment to a certain degree. If the APs are not initially located correctly in your environment AirMatch may not be able to overcome poor physical deployment of APs.

ClientMatch

There is an eminent need for directing the clients to the best suited APs based on dynamic environment variables in order to achieve the best network performance. ClientMatch constantly optimizes the client association by continuously scanning the wireless environment and sharing information about the clients and the APs. Based on the dynamic data obtained, the clients are steered to the most suitable AP. No software changes are required in the clients to achieve this functionality.

ClientMatch looks at the client’s view of the network while making a steering decision. The client view of the network is built with all the probe requests that are received from the same client on different APs. This is used to build the Virtual Beacon Report which forms the building block for ClientMatch which helps clients that tend to stay associated to an AP despite deteriorating signal levels. ClientMatch

will continuously monitor the client's RSSI while it is associated to an AP, and if needed, will try to move the client to a radio that would give it a better experience. This prevents mobile clients from remaining associated to an AP with less-than-ideal RSSI, which can cause poor connectivity and reduce performance for other clients associated with the AP.

ClientMatch continually monitors the RF neighborhood for each client to provide ongoing client band-steering and load balancing, and enhanced AP reassignment for roaming mobile clients. Since the client is not aware of AP load from a client count and channel utilization perspective, ClientMatch will help the client make a better decision as to which AP it is connected to. ClientMatch is not a roaming assistant, but rather a network optimization tool to enhance a client's decision when selecting the proper AP.

ClientMatch is Wi-Fi 6 aware and there is no special knob for this feature because it is enabled by default but can be disabled if required. It will try to match Wi-Fi 6 clients to Wi-Fi 6 radios in a mixed AP deployment environment.

Roaming Optimizations

Each AP deployment and its RF environment are unique, so the best practice for one environment may not be the same for another. Still, certain components will help clients for seamless roaming and achieve an overall better user experience. Wireless client devices are extremely sensitive to RF environment. Device performance can be substantially improved by using the following recommendations.

Table 2: Roaming recommendations

Feature	Recommendation	Description
Transmit Power	Leave it at default values and let AirMatch take care of these values.	Adjusting the AP's transmit power using Aruba's AirMatch technology is recommended for optimal roaming performance. Leave it at default values (5 GHz: Min 18 / Max 21 dBm and 2.4 GHz: Min 6 / Max 12 dBm) and let AirMatch take care of this configuration. In case manual configuration is required then please note that the difference between minimum and maximum Tx power on the same radio should not be more than 6 dBm. Tx power of 5 GHz radio should be 6 dBm higher than that of 2.4 GHz radio. Having the radio power on the 2.4 GHz radio lower allows for the 2.4 GHz band to be less attractive to a client and influencing the use of a 5 GHz radio. Since 2.4 GHz is roughly twice as strong of a signal, a lower dBm power is needed to make the radio less attractive to a client. Some clients that are dual band capable will still prefer a 2.4 GHz radio over a 5 GHz radio if the available power for both is the same. The best solution will drive as many clients to a 5 GHz radio as possible but allow 2.4 GHz-only clients to maintain connectivity. Setting a consistent power levels across all available radios leads to more predictable roaming behavior among a group of APs.
Channel width	Let AirMatch decide the optimal channel and the channel width suited for the particular RF environment	Use 80 MHz channel with DFS channels only if no radar signal interference is detected near your facility. Enabling DFS channels could create coverage holes for clients who do not support it. Most of the clients do not scan DFS channels initially; this makes roaming more inconsistent when using these channels.
Band Steering	Enable 11ax aware Client Match	ClientMatch optimizes user experience by steering clients to the best AP w.r.t client capabilities and AP load.
Local probe request threshold	15	Prevents APs from responding to a client's probe request if their signal to noise ratio value is below 15 dB, thereby encouraging roaming to closer APs which are responding to the client's probe request.

Table 3: Fast roaming recommendations

Feature	Recommendation	Description
Opportunistic Key Caching (OKC)	Enable	Avoids full dot1x key exchange during roaming by caching the opportunistic key. **NOTE:** MacOS and iOS devices do not support OKC.
802.11r Fast BSS Transition	Enable	802.11r enables clients to roam faster and recent macOS, iOS, Android clients, and Win10 are support the protocol. Some older 802.11n devices, handheld scanners and printers may have connectivity issues with 802.11r enabled on WLAN. This feature is disabled by default.
802.11k	Enable	Enable 802.11k with the following changes. Set the Beacon Report to Active Channel Report and disable the Quiet Information Element parameter from the Radio Resource Management profile.

Air Slice

Air Slice is a unique RF feature that leverages Wi-Fi 6 core constructs to optimize user and application experience. By combining Aruba's stateful firewall and Layer 7 Deep Packet Inspection (DPI) to identify user roles and applications, the APs will dynamically allocate the bandwidth and other RF resources that are required to meet the performance metrics for business-critical applications to ensure better user experience. Using Air Slice, the network administrator can further orchestrate radio resources to work with ClientMatch to go beyond the traditional capabilities of Airtime Fairness.

Air Slice uses internal hardware queues to prioritize traffic within the same access class. Zoom video can be prioritized over other video traffic with the same WMM tag which means it crosses the barriers of traditional WMM QoS. WMM boost is also implemented to auto increase WMM priority for applications that do not have DSCP/ToS markings set, which shows how best-effort traffic can be prioritized using Air Slice. Air Slice provides benefit for non-11ax clients by using adaptive priority queuing for the enterprise application flows and WMM boost.

A growing number of enterprises are using latency-sensitive, bandwidth-demanding applications like Augmented Reality or Virtual Reality, or other collaborative applications such as Zoom, Skype for Business, and Slack. These applications have stringent quality of service requirements. If these QoS parameters are not met, it translates into poor user experience; moreover, IoT devices such as security cameras or HVAC sensors are also becoming prevalent, and their requirements are very different in terms of sleep cycles and latency sensitivity.

Since many of these new applications have stringent QoS requirements in terms of latency, bandwidth, and throughput, an enhanced QoS is needed. Air Slice should be enabled in order to improve the user experience while using latency sensitive applications.

The following table lists the applications supported by default with Air Slice.

Table 4: Air Slice Default Applications

Default Applications	
Wi-Fi Calling	Zoom
Office 365	Skype For Business
GoToMeeting	Slack
Cisco WebEx	Amazon Web Services
Dropbox	GitHub
Custom Applications (Up to 5)	

Access Point Placement

Aruba recommends doing a site survey for all wireless network installations, whether you use a virtual tool or physical site survey for special cases like manufacturing floors and warehouses. 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 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. A key thing to remember for AP placement is that RF signals with higher frequency cover short distance compared to the low-frequency signals. Therefore, APs should be placed in such a way that the 5 GHz band signal covers the target area.

For Wi-Fi 6 AP deployment, the minimum Received Signal Strength Indicator (RSSI), which is a measurement of how well client device can hear a signal from an access point, should be -55 dBm throughout the coverage area. The reason for choosing an RSSI better than -55 is so the APs can reliably provide a Modulation and Coding Scheme (MCS) with at least an MCS11 data rate on 40 MHz high density deployments. MCS rates dictate both the technology chosen and the transmit and receive rates for the wireless client. Wi-Fi 6 clients with a poor MCS rate will roll back to older technologies like Wi-Fi 5 and use a lower transmit and receive rate for successful data transmission.

While deploying a Wi-Fi 6 network using dual-band APs, 2.4 GHz radios of some of the APs should be turned off to avoid co-channel interference. Sufficient coverage should be validated using heat maps to make sure there are no coverage holes.

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 layout with APs. The staggered spacing between APs is equal in all directions and ensures suitable coverage with seamless roaming.

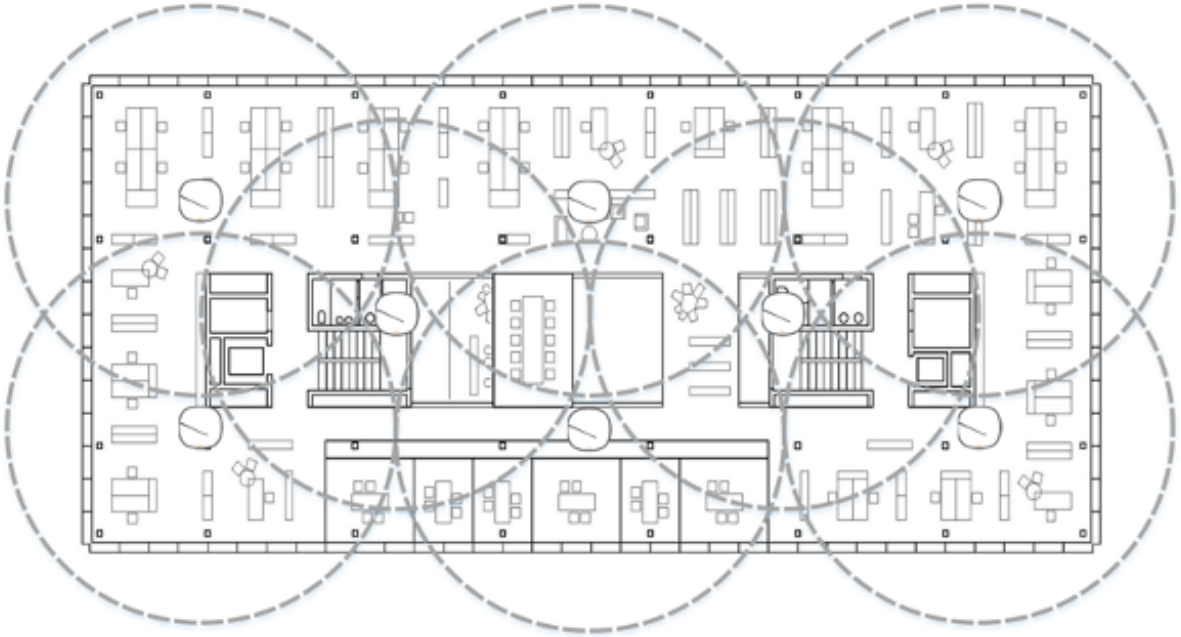

Figure 24: Sample office AP layout

After studying an environment with the 35-50-foot (10-15 meter) rule in mind, make sure there is 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. If there are high-density areas such as large conference rooms, cafeterias, or auditoriums, additional APs may be needed.

AP Mounting Recommendations

Indoor APs are typically deployed in ceiling or wall mount deployments. Aruba does not recommend desk or cubical mounted deployments. These locations typically do not allow for a clear line-of-sight throughout the coverage area, which in turn can reduce overall WLAN performance. With the exception of the hospitality style AP, APs with internal antennae should not be wall mounted if at all possible. Wall mounted APs should have external antennae and Aruba has a selection of antennae options available. There are some cases where wall mounting an internal antennae AP on the wall is valid however this kind of design should include professional services to validate proper roaming and coverage patterns base on the AP model selected.

Channel Planning

The Aruba AirMatch software will handle automating channel and power assignment for even the most challenging wireless installations. If you want to plan your channels following the details in this section, please contact an Aruba or partner systems engineer or consulting systems engineer for verification of the 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 for North America 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. Aruba recommends using only these three channels for 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 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.

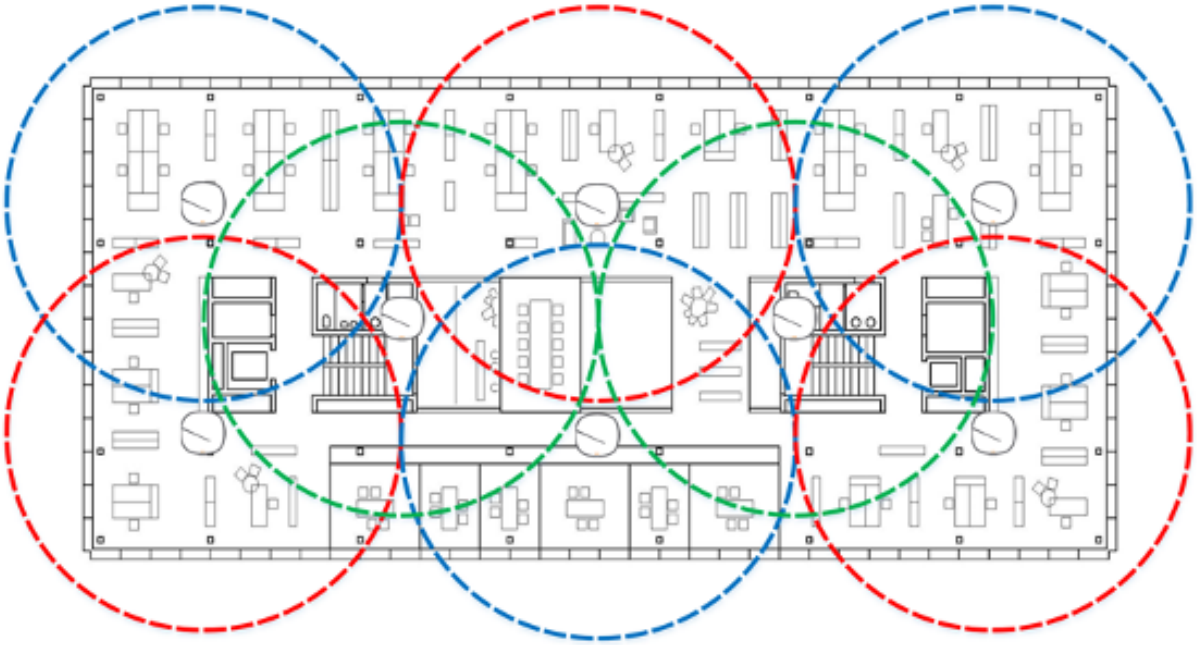

Figure 25: Channel layout for 2.4 GHz band with three unique channels

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. With more channels in the 5 GHz band, Aruba recommends all capable clients connect on 5 GHz and 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 result 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 there is a high-density open-office environment or a loss of channels due to DFS interference, it is better to start with 20 MHz channels. Dynamic Frequency Selection (DFS) is a Wi-Fi function that enables WLANs to use 5 GHz frequencies that are generally reserved for radars.

Due to the high number of APs and increasing number of connected devices, there are certain office environments that would benefit from 80 MHz-wide channels and the much wider 160 MHz channels. The wider channels will make sense once there are enough Wi-Fi 6 clients to take advantage of the new features outlined in the Wi-Fi 6 Enhancement section below. The following figure highlights the 40 MHz channel allocation for the 5 GHz band and shows the DFS channels.

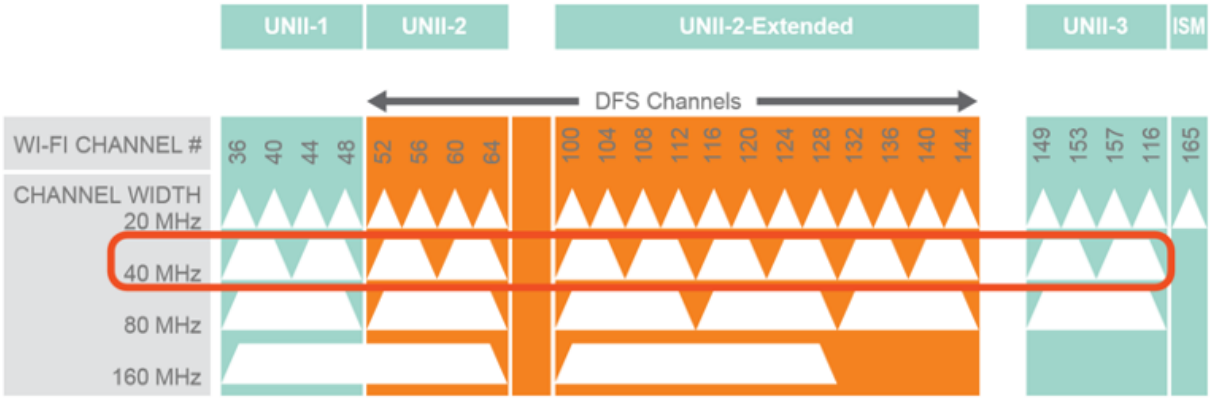

Figure 26: Channel allocation for the 5 GHz band

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. If an AP detects radar transmissions on a channel, the AP stops transmitting on that channel for a time and moves to another channel. It is recommended to use AirMatch for channel allocation as it will detect interference and accommodate the optimal channel plan to avoid active DFS channels.

In the past, it was common to disable all DFS channels, but today most organizations attempt to use the channels available in their country. In some areas, DFS channels overlap with radar systems. 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 40 MHz-wide channels, there are up to 12 channels available. Depending on local regulations and interference from radar or other outside sources, the total number of usable channels will vary from location to location. To get maximum performance, it is recommended to allow AirMatch to automatically determine the channel width whenever possible. It is also recommended not to disable any channel widths in case AirMatch determines they will work in the environment.

NOTE:

You can find a list of the 5 GHz channels available in different countries at the following: 5 GHz WLAN Channels by Country

Power Settings

The optimum power settings vary based on the physical environment and the initial settings should always follow the recommendations from an early site survey. For the long term, Aruba recommends using AirMatch to decide the optimal transmit power values for each AP. AirMatch uses telemetry data from the entire network to compute transmit power values unique for the particular deployment. Please refer to the AirMatch section for more details.

When not using AirMatch, the following guidelines for a typical wireless design are recommended: - In the 2.4 GHz band, set the minimum power threshold to 6 dBm and the maximum power to 12 dBm - In the 5 GHz band, set the minimum power threshold to 18 dBm and the maximum to 21 dBm - Do not exceed a power level difference of 6 dBm between the minimum and maximum settings on all radio bands - The minimum power level differences between equal coverage level 2.4 GHz radios and 5 GHz radios should be 6 dBm

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 use, the better overall throughput is for a higher number of devices. As mentioned previously, a typical Aruba wireless installation uses the AirMatch software and AI Insights running in the cloud for RF channel planning and optimization.

Proxy ARP on Gateway

Enabling this feature will inform the Gateway to ARP on behalf of a client in the user table. When enabled on a VLAN with an IP address, the Gateway will provide its MAC address in the Proxy ARP. If the VLAN does not have an IP address the Gateway will supply the MAC address of the client in the user table. This feature is off by default and should only be changed to address specific deployment scenarios where the Gateway is a transparent hop to another device, for example when using Aruba VIA VPN.

NAT/Routing on Gateway

Campus installations of a Gateway should always be Layer 2 and the Gateway should not perform Layer 3 operations. The client's default gateway should be another device like a router or switch and the Layer 2 network should be dedicated for the clients attached to the Gateway. The broadcast and multicast management features of the Gateway allows large subnets to be used without issue. It is recommended the Layer 2 network be as large as supported by the Gateway and the supporting switching infrastructure. Table sizes, ARP learning rates, physical layer rates, and redundancy all are factors to account for in the switching infrastructure.

NOTE:

Firewall policies must be used when routing is enabled on the Gateway to control inter-VLAN traffic and determine whether traffic should be routed or NAT'd.

Network Resiliency

Aruba's recommended network design is a highly available, fault tolerant network. There are features that should be enabled from a software perspective to ensure the network is prepared for interruptions. This section will provide general guidelines for software solutions that provide fault tolerance and allow for upgrades with minimal service impact

Wireless Resiliency Technologies

For campus wireless Aruba recommends either APs on their own or APs with Gateways. In both of these designs there are features that can be enabled to ensure that the network is highly resilient.

Authentication State/Key Sync

Authentication keys are synchronized across APs by the Key Management Service (KMS) in Central. This allows a client to roam between APs without re-authenticating or rekeying their encrypted traffic. This decreases the load on the RADIUS servers, but also speeds up the roaming process for a seamless user experience. Key synchronization and management are automatically handled by the APs and Central, so no additional user configuration is required.

Firewall State Sync

Traffic from a client can be synchronized across a primary and secondary Gateway when using a cluster. This allows the client to seamlessly fail from its primary Gateway to a secondary Gateway. The system synchronizes encryption keys between Access Points so that when a client moves to its secondary Gateway the client does not need to re-auth or rekey its encrypted traffic. These two operations working together make it completely transparent to the client when they move between Gateways or APs. This is a key component to Aruba's high availability design and Live Upgrade features. When using a bridged SSID the firewall state is synced upon a roaming event from a client for a seamless roaming event with no traffic disruption.

Cluster Design Failure Domain

If more than one Gateway fails, there is a possibility of an outage to some clients. When a Gateway fails, the clients with a single Gateway connection will be recalculated, and the cluster will be rebalanced. This doesn't happen immediately and there isn't a set amount of time as it depends on the number of users. If a second Gateway fails before the rebalancing can occur the client will be disassociated and reconnect to an available Gateway. This will look like a dirty roam, but the client will reestablish a connection as long as the Gateways are not at capacity during a failure. To mitigate a multiple Gateway failure, a resilient deployment should minimize the common points of failure. Using disparate line cards or switches, multiple uplinks spanning line cards or switches, port configuration validation, and multiple Gateways are all foundational requirements to limit the failure domain.

Switching Resiliency Technologies

When it comes to the campus switches, Aruba recommends either a two-tier LAN with collapsed core or a three-tier LAN with a routed core. In both of these designs there are common features that can be enabled to ensure that the network is highly resilient. The two-tier campus is on the left and the three-tier campus is on the right.

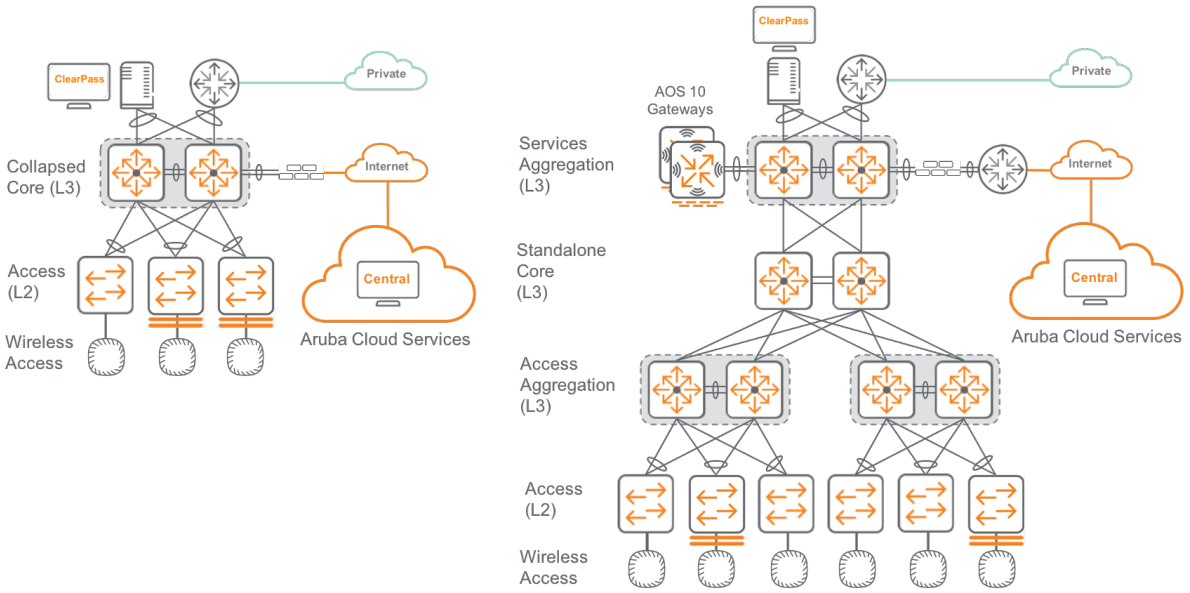

Figure 27: Two-tier and three-tier wired

Virtual switching Extension

VSX is a virtualization technology used to logically combine two AOS-CX switches into a single logical device. From a management/control plane perspective, each switch is independent of the other, while the Layer 2 switch ports are treated like a single logical switch. VSX is supported on 6400, 8320, 8325, or 8400 models, but it is not supported on Aruba CX 6300, 6200 or 6100 models. VSX should only be enabled if the devices are positioned in a collapsed core or aggregation layer. It is important to note that a VSX pair cannot be mixed between different models meaning an 8325 cannot be mixed with 8320. Here is the list of supported combinations:

- Aruba CX 6400: All combinations within the 6400 family are supported
- Aruba CX 8320: All combinations within the 8320 family are supported
- Aruba CX 8325: All combinations within the 8325 family are supported
- Aruba CX 8400: All combinations within the 8400 family are supported

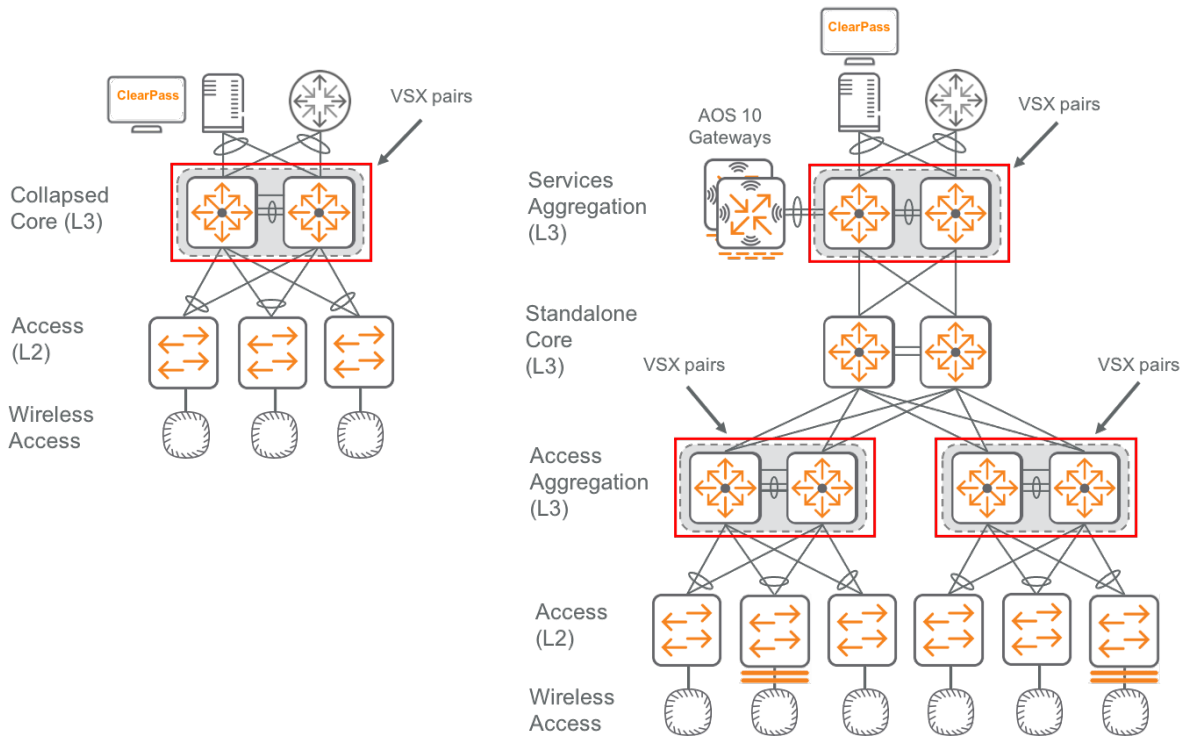

Figure 28: VSX pair placement

VSX pairs stay in sync using synchronization software and hardware tables shared over the inter-switch link (ISL). The ISL is a standard link aggregation group (LAG) designated to run the ISL protocol between the paired devices. VSX allows the two devices to appear as a single device using an Active Gateway feature which is a shared IP and MAC address. Each VSX pair appears as a single Layer 2 switch to a common downstream access switch utilizing a specialized LAG called a MC-LAG. Multi-Chassis Link Aggregation MC-LAG allows the aggregation layer switch pair to appear as a single device to other devices in the network, such as the dual-connected 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.

The access switch uses a standard LAG connection and from the access switch perspective, the VSX pair appear as a single upstream switch. This minimizes the fault domains for links by separating the connections between the two VSX paired switches. This also minimizes the service impact with the live upgrade feature because each device has its own independent control plane and link to the downstream access devices.

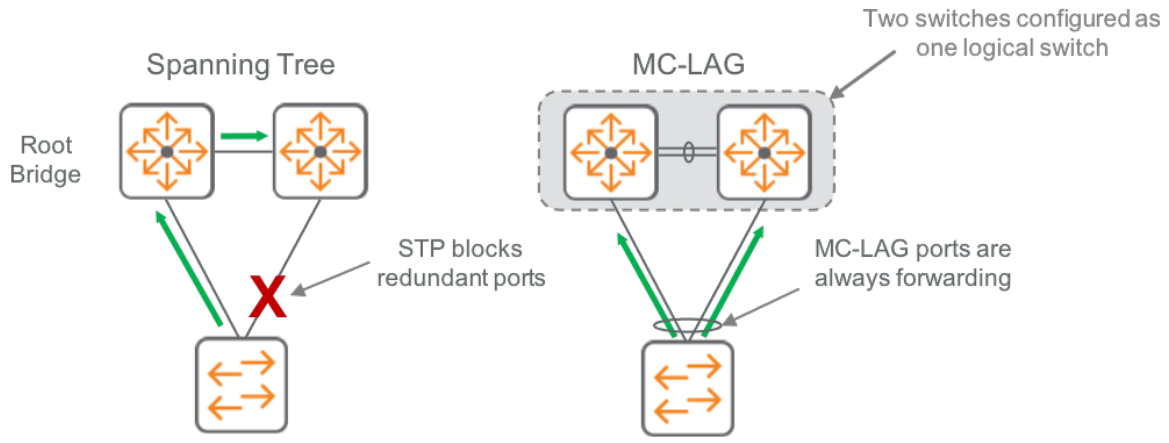

Figure 29: Traditional spanning tree vs VSX

NOTE:

When using LAG or MC-LAG, STP is not required, but should be enabled as an additional enhanced loop protection security mechanism.

LACP - The link aggregation control protocol (LACP) combines two or more physical ports into a single trunk interface for redundancy and increased capacity.

LACP Fallback - LAGs with LACP Fallback enabled allow an active LACP interface to establish a connection with its peer before it receives LACP PDUs. This feature is useful for access switches using Zero Touch Provisioning (ZTP) connecting to LACP configured aggregation switches.

Inter-switch Link - The best practice for configuring the ISL LAG is to permit all VLANs. Specifying a restrictive list of VLANs is valid if the network administrator wants more control.

MC-LAG - These LAGs should be configured with the specific VLANs and use LACP active mode. MC-LAGs should NOT be configured with a permit all VLAN.

VSX Keepalive - The VSX keepalive is a UDP probe which sends hellos between the two VSX nodes and is used to detect a split-brain situation. The keepalives should be enabled with a direct IP connection between the VSX pairs in a designated VRF domain. The VSX keepalive is not yet supported over the Out-of-Band Management (OOBM) port.

Active-Gateway - This is the default-gateway for endpoints within the subnet and it needs to be configured on both VSX primary and secondary switches. Both devices must also have the same virtual MAC address configured from the Private MAC address spaces listed below. There are four ranges reserved for private use.

- x2-xx-xx-xx-xx-xx
- x6-xx-xx-xx-xx-xx
- xA-xx-xx-xx-xx-xx
- xE-xx-xx-xx-xx-xx

NOTE:

x is any Hexadecimal value.

PIM Dual - If the network is multicast enabled, the default PIM Designated Router is the VSX primary switch. In order to avoid a long convergence time in case of a VSX primary failure, the VSX secondary can also establish PIM peering. Aruba recommends configuring PIM active-active for each VSX pair.

VSX and Rapid PVST - Aruba recommends configuring VSX and spanning tree per the design guidelines outlined in the document below.

NOTE:

Certain VSX use cases fall outside of the design guidance in this document, but they are covered in detail in the VSX Configuration Best Practices guide. The best practices guide includes in-depth information about spanning tree interactions, traffic flows, active forwarding and the live upgrade process.

Virtual switching Framework

Stacking allows multiple access switches connected to each other and 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.

AOS-CX access switches provides front plane stacking using the Virtual switching Framework (VSF) feature, utilizing two of the four front panel SFP ports operating at 10G, 25G, or 50G speeds. VSF Combines the control and management plane of both switches in a VSF stack which allows for simpler management and redundancy in the access closet. VSF is supported on the 6300 and 6200 model of switches.

VSF supports up to 10 members on a 6300 an up to 8 members on a 6200, Aruba recommends a ring topology for the stacked switches. A ring topology, which can be used for 2 switches all the way up to 10 switches, allows for a fault tolerance in the case of link failure because the devices can still reach the commander or standby switch using the secondary path. The commander and standby switch should have separate connections to the pair of upstream aggregation switches. If the commander fails, the standby switch can still forward traffic upstream minimizing the failure domain to just the commander switch. The recommended interface for switch stacking links is a 50G Direct Attach Cable (DAC) which will allow enough bandwidth for traffic across members.

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.

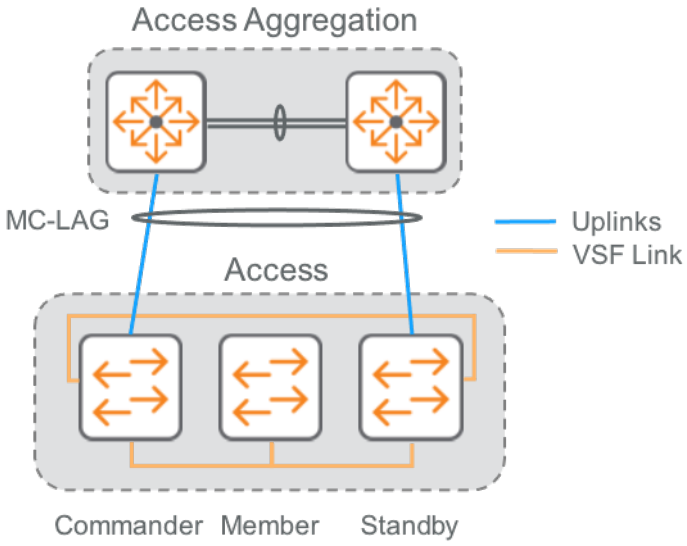

Figure 30: VSF connections

In order to mitigate the effects of a VSF split stack, a split-detection mechanism must be enabled for the commander and standby members of the stack which is known as Multi-Active Detection (MAD). This is done using a connection between the OOBM ports on the primary and secondary members to detect when a split has occurred. Aruba recommends the OOBM ports are directly connected using an Ethernet cable.

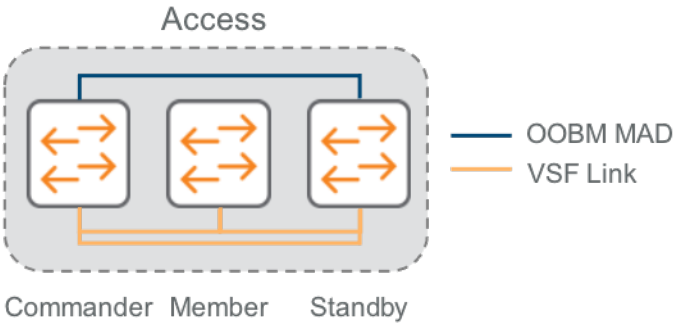

Figure 31: VSF OOBM MAD and links

Campus Policy Layer

The *Policy Layer* runs on top of the connectivity layer and is responsible for the creation, maintenance and enforcement of network security policies. An endpoint security posture is determined from information gathered using multiple points of context such as date and time, geographic location, authentication method, device type, operating system, as well as other behavioral aspects of the device. After a determination is made, a user role is assigned, and policy can be applied based on the trust level of the device. The ESP architecture supports consistent policy for wired and wireless that includes options for access control at the edge of the network and also traffic tunneling for additional inspection by the stateful firewall on the Gateway.

Endpoint device policy is a pivotal component of the Aruba ESP architecture because it gives an organization a way to create business logic in the network. From the business perspective, this allows dynamic connectivity, a consistent experience across wired and wireless networks, as well as increased network security. From a technical perspective, devices are identified using RADIUS and device profiles. Once an endpoint device has been identified and a user role assigned, the network can drive the security posture. Aruba supports micro and macro segmentation which, depending on the use case, can be used separately or together. The Aruba ESP Architecture achieves this using Dynamic Segmentation, which covers a range of features that include User-Based Tunneling, Virtual Network-Based Tunneling, Device Profiles and User Roles. This section will cover the different ways policy can be implemented within the ESP campus architecture.

Network Authentication

The main responsibility of a secure network is to guarantee that only authorized users are allowed access to the network. Security protocols need to let authorized users in, while keeping the rest out. There are several ways to authenticate and authorize users but using a centralized service to maintain user and device profiles is the easiest approach. A central service running a secure networking protocol provides better security and ease of maintenance, allowing an organization to apply policy from a single administered network point, which also makes it simpler to track usage.

Local Accounts

Outside of the default accounts, local accounts should not be created. All management access should be authenticated through TACACS or RADIUS for accountability purposes. Service accounts are included in this, the local admin/root account should not be used for service accounts or monitoring systems. Any local account that is deemed a necessity should comply with password complexity requirements and password rotation requirements set by the organization. Account information on any local account should also be limited to as few people as possible with strict access and usage requirements.

Session Timeouts

Users that have connected to Central, Gateways or switches should not be allowed to stay connected for extended periods of time without any activity. This is to protect the number of open sessions to the devices as well as anyone who could potentially edit something on the devices that should not have access via an idle connection. The Recommended Timeout can vary between 1 and 5 minutes depending on security policy of the environment. When setting the session timeout ensure that it is consistent across the entire network.

Minimum password

Complex passwords are a requirement for just about everything these days and there is no exception for the network infrastructure. Having complex passwords drastically decrease the likelihood of a password cracking. It is best practice across all Aruba platforms to have a complex password. Passwords should be 8 characters or more, with a combination of capital letters, lower case letters, numbers and special characters. It is also recommended that few to no local accounts exist on devices and the admin/root account password is set to a 64-character complex password and stored in a safe or password management app. Setting this long of a password and storing it in a safe will typically mitigate the need to rotate the local account on a regular basis. All management access should be authenticated through TACACS or RADIUS for accountability purposes. Service accounts are included in this, the local admin/root account should not be used for service accounts or monitoring systems. With RADIUS and TACACS the password complexity is typically controlled by a directory service that will comply with password age and complexity requirements set by the organization.

SSH

SSH is one of the most common ways to connect to Network infrastructure as it is the most secure way to interact with network devices. It is important to ensure that the strongest ciphers are being used to connect into the network devices. Pending the region and supported ciphers on different systems it is recommended to review the Aruba hardening guide for details on which cipher should be selected. The Hardening guide can be found on the Aruba support site.

TACACS

Terminal Access Control System (TACACS) is a security protocol that provides centralized authentication and validation of users who are attempting to access to a network device. TACACS servers should be integrated with an existing directory service to enforce group permission, password complexity, and password life cycle requirements. It is recommended to have all Aruba devices TACACS enabled for authentication, command authorization, and auditing. Aruba ClearPass Policy manager supports TACACS services for any hardware vendor that supports the TACACS protocol including Aruba devices.

Certificate Replacement

Replacing the factory certificates is part of the Aruba hardening guide. Certificates should always be replaced before being devices go into production. TLS 1.2 should be the only Cipher enabled on any web interface. On Gateways you also have the ability to set the Cipher strength to high which is recommended as part of the hardening guide. Using a wildcard certificate for web interfaces for management is acceptable across all Aruba platforms.

API Access

When using any API interface, it is critical that the interface is protected. If the API interface is enabled, it is recommended an ACL be enabled so that only a management subnet is able to access that interface. Depending on the Aruba product the management interface may no longer be required to be enabled. If devices are managed through Aruba Central there is no longer a need to enable the local web interface and it can be disabled.

RADIUS

Remote Authentication Dial-In User Service (RADIUS) is a secure networking protocol to enforce security policies on user devices connecting to the network. The network infrastructure becomes a client to the RADIUS server and a user device connecting to the network is authenticated by the RADIUS server which permits or denies the user device. When using a RADIUS sever, the network will no longer have control over what happens to a user device connected to an enabled interface. Besides just the users that should be authenticated, it is important consider the authentication priority, Change of Authorization (CoA), Critical Authentication and Rejection Roles.

Authentication Priority - There are two main authentication methods: 802.1X and MAC authentication. It is recommended to have access switches and Gateways prefer 802.1X and then, fall back to MAC authentication.

Change of Authorization (CoA) - After a device has been authenticated, the authorization can be modified through a CoA. This is typically done for devices that have triggered a security alert and need to be moved to a quarantine VLAN. It is recommended to enable CoA for all switches and Gateways.

Critical Authentication - If the RADIUS sever fails or authentication request are not getting to the server, it is important to have a critical authentication role on the switching infrastructure to allow devices on the network while the RADIUS server is unreachable. The critical authentication role by itself will not allow basic network connectivity, so it is important to ensure that there is a role with network connectivity to critical applications and Internet access. Remember, devices in the critical authentication role may not be trusted, so their access should only allow basic connectivity to prevent potential exploitations.

Rejection Role - When devices do not have the appropriate credentials to connect to the network, the rejection role should restrict, but not deny all access. The rejection role should place the device in a VLAN where there is basic internet connectivity, but also allow the device to be re-authenticated and placed in the correct VLAN.

Client limit - This will limit the number of clients allowed to on a port, this will also protect users from a potential mac flooding attack as the number of clients cannot go past this limit. The maximum number of clients on a port should be 5, however this is flexible depending on the deployment.

User Role

The benefits of assigning roles to user and devices is well known within the Aruba wireless portfolio. This concept has been adapted to the Aruba wired portfolio and eases the burden of configuration by grouping policies into a role that can be referenced by the Administrator. When ClearPass is used, time of day, type of machine, and device profiling are available when deciding if a user and their device should be allowed on the network, and what access rights are granted.

The following is a list of key points regarding User Roles:

- A User Role will dictate if traffic is locally switched or tunneled back to the Gateway.
- Any RADIUS server can be used to return a Local User Role (LUR).
- LURs are assigned from the RADIUS server using the 'HPE-User-Role' VSA.

A role will dictate what VLAN is assigned, tagged or untagged, and if the traffic is locally switched or tunneled back to a Gateway. Optionally, a role can also assign a policy ACL or QOS, re-authentication timers, and captive portal redirect. It doesn't matter if the role is pre-defined or downloaded from ClearPass, it must exist on the network device before it can be applied.

Dynamic Segmentation

Dynamic Segmentation unifies policy enforcement across wired and wireless networks, keeping traffic secure and separate. It utilizes intelligence gathered from Aruba's role-based policy capability, user firewalls, Layer 7 application visibility and integrated web content filtering. Dynamic Segmentation has two main components. The first one assigns policy or User Roles dynamically to a device based on authentication. This means network administrators do not have to preconfigure access switch ports or APs with VLANs for each device that is connected. This allows an organization to use colorless ports on the wired side which significantly reduces the operations planning for initial deployments and on-going changes to accommodate moves, adds, and changes.

The second component is the segmentation of traffic on a given subnet using local VLANs or by tunneling traffic back to centralized Gateways for further inspection and micro-segmentation. The Gateway's advanced firewall feature set provides enhanced security and performance benefits based on increased visibility from profiling and deep packet inspection of traffic.

Untrusted traffic on the network from wired and wireless devices is tunneled to centralized Gateways. They inspect the traffic and apply security policy before allowing it to continue. Traffic can be micro-segmented to prevent east-west flows within a given subnet or it can be prevented from traveling north-south to other locations within the campus.

User-Based Tunneling

On the wired side, User-Based Tunneling (UBT) sends device traffic to centralized Gateways, which apply a consistent set of rules with a User Role. UBT provides the flexibility to maintain existing VLAN segmentation which allows for simple adoption and an easy migration path. Traffic added to the policy overlay can coexist with traffic using the connectivity underlay without disruption of services.

Policy is enforced at the Gateway using intelligence gathered from a User Role defined in ClearPass combined with the Layer 7 firewall capability. Dynamic Segmentation provides unique context on user attributes, ranging from their role in the organization, type of device, and their network location. This helps an organization deploy a Zero Trust security policy in their environment by providing a consistent set of rules for wired and wireless users.

A User Role can contain an ACL policy, QoS policy, captive portal redirection, VLAN assignment, and device attributes. When the User Role is received from the RADIUS server, a command to redirect traffic to a Gateway can be included. When traffic is redirected to a Gateway, the authentication sub-system provides a secondary role to the Gateway. The secondary role is the User Role on the Gateway where policy exists for tunneled users applying the advanced firewall and security rules. This secondary role information is an indication to the Gateway that it has to enforce additional policies to the traffic based on the configuration associated with the secondary role and then, form the tunnel. The tunnel also provides high availability and load balancing with the Gateway clustering feature to extend resilient secure access to all wired devices within the Intelligent Edge network.

The following diagram shows UBT in the three-tier wired design. The employee, visitor and IoT devices have different access capabilities based on their specific User Role policies even though they come from the same VLAN subnet in the access layer.

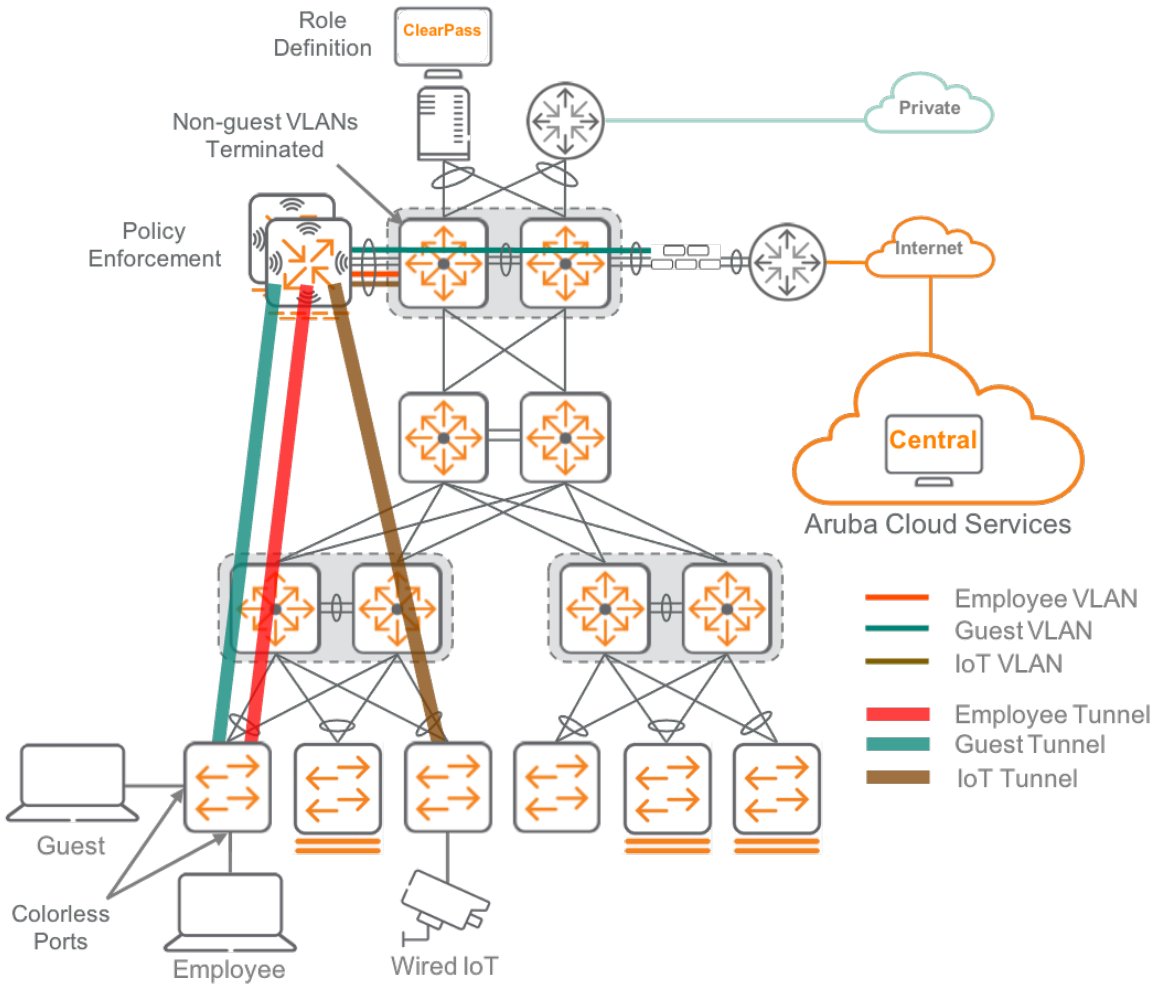

Figure 32: **Policy enforcement using UBT in three-tier wired**

Colorless Ports

Colorless ports are a wired solution that do not require static configurations of VLANs to access ports. They allow the administrator to configure access switches generically with only a management VLAN on each port. When a device is connected, a policy is automatically applied to the connected device using a local or downloadable user role in real-time.

The following is a list of key points regarding colorless ports on Aruba switches:

- Colorless Ports are the “dynamic” in Dynamic Segmentation because there is no need to assign user VLANs to access-ports.
- Colorless Ports are not exclusive to user-based tunneling, nor do they require ClearPass.
- Colorless Ports can be used with LURs.

Tunnel-mode SSID

On the wireless side, tunnel-mode SSIDs consist of APs with Layer 2 tunnels to centralized Gateways, which apply the same set of rules for all wired and wireless devices. The policy is configured in Central or ClearPass and the Gateway acts as the policy enforcement point with intelligence gathered from the role-based policy and firewall capabilities. The tunnel-mode SSIDs are terminated at the Gateway which offer micro-segmentation for all traffic using the stateful firewall. Dynamic Segmentation provides unique context on user attributes, ranging from their role in the organization, type of device, and their network location. This helps an organization deploy a Zero Trust Security policy in their environment by providing a consistent set of rules for wired and wireless users. From the scaling perspective, Gateways are recommended when customers expect to deploy more than 500 APs and 5000 clients at a single site.

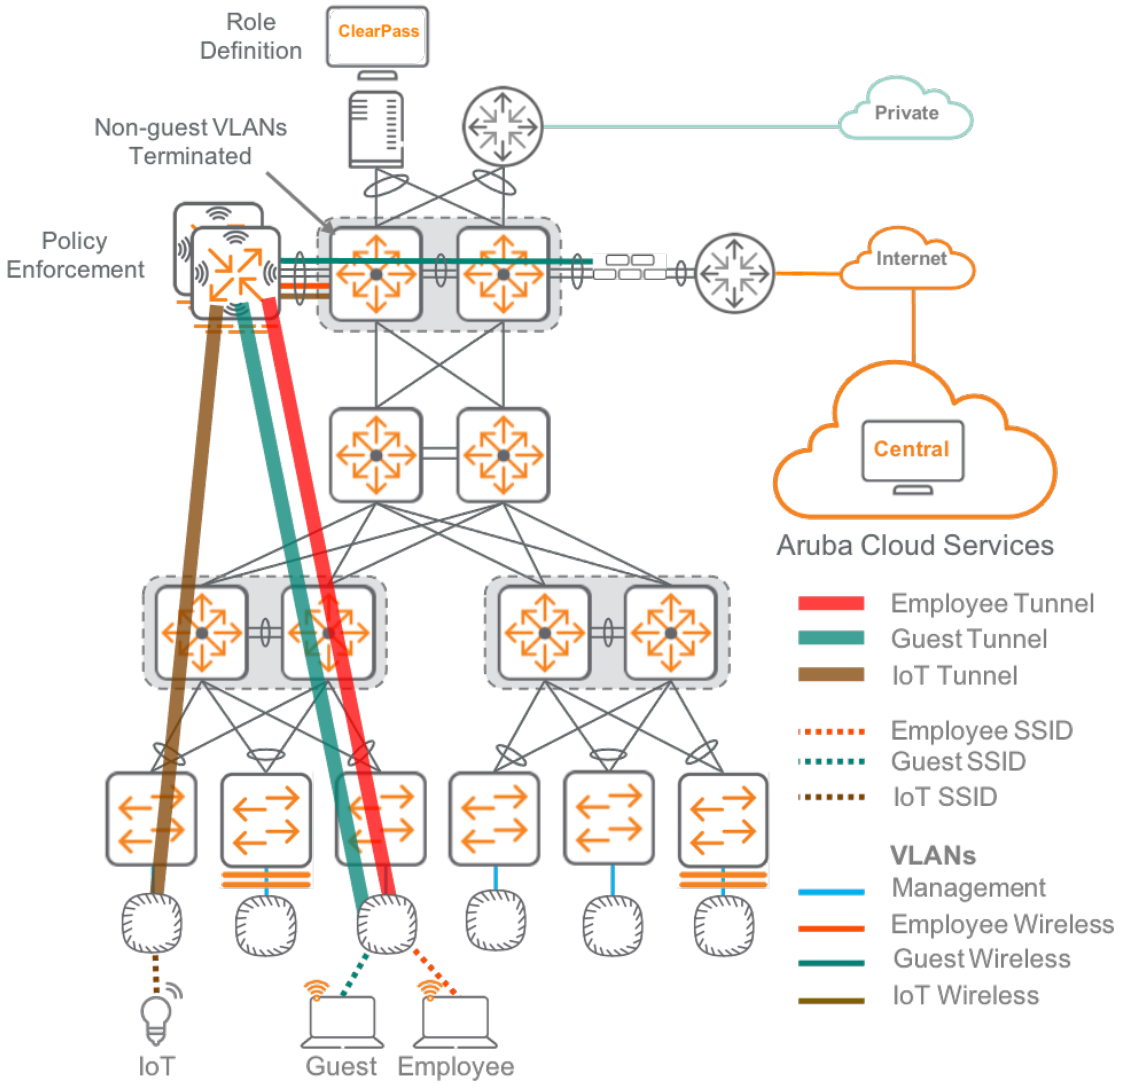

Figure 33: Tunnel-mode in three-tier wired

Network Segmentation

Trusted traffic on the network can be controlled at the switch or AP when an organization does not want to tunnel to a centralized location for inspection. Wired traffic will use LUR assigned VLANs and Layer 3 / Layer 4 ACLs to prevent north-south access to other areas of the network. Wireless traffic will use bridge-mode SSIDs to segment traffic into the existing VLANs on the access switches.

Access Control Lists

ACLs identify traffic based on IP address, port numbers or both. Once the traffic has been identified, ACL's have the option to permit, deny, tag or log. ACLs can be used to deny traffic from one destination to another, to identify traffic for QoS or to filter routes in a route map. In addition to the three instances, there are three types of ACL's supported on Aruba devices, MAC-based, IPv4 and IPv6 ACL's. This section will not cover QoS classification, because it is covered in the "Quality of Service" section of the guide. Route maps are used when redistributing traffic from one routing instance to another or in the case of route leaking between VRFs. Guidance for how to write route map rules is specific to each network installation.

When using ACLs to segment traffic, it is best to use an IPv4 ACL and have distinguishable network boundaries. If an ACL is used to segment devices, it is important these devices are not within the same subnet because traffic within a subnet does not pass through the Layer 3 portion of the switch where the ACL is applied. The recommended ACL denies user VLANs (BYOD, Employee and Visitor) from accessing the network infrastructure management VLAN with SSH or HTTP/HTTPS.

On AOS-CX devices, ACLs can only be applied to an interface in specific direction depending on the platform, so it is important to note which direction ACLs are applied in order to have a desired result. The following tables list the AOS-CX capability to apply ACL inbound or outbound per interface type and per ACL type.

Table 5: AOS-CX ACL inbound and outbound per interface and ACL type

ACL Interface Type	8400X	8360	8325	8320	64xx	6300	6200	6100
Ingress v4 ACL on ports	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ingress v4 ACL on vlans	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ingress Routed v4 ACL on vlans	Yes	Yes	Yes	Yes	Yes	Yes	-	-
Ingress v6 ACL on ports	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ingress v6 ACL on vlans	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ingress Routed v6 ACL on vlans	Yes	Yes	Yes	Yes	Yes	Yes	-	-
Ingress MAC ACL on ports	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ingress MAC ACL on vlans	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Egress v4 ACL (on route-only ports)	Yes	Yes	Yes	Yes	Yes	Yes	-	-

ACL Interface Type	8400X	8360	8325	8320	64xx	6300	6200	6100
Egress v4 ACL (on bridged ports)	-	Yes	-	-	Yes	Yes	Yes	-
Egress Routed v4 ACL on vlans	-	Yes	Yes	Yes	Yes	Yes	-	-
Egress v6 ACL on ports	-	Yes	-	-	Yes	Yes	Yes	-
Egress v4 ACL on vlans	-	Yes	Yes	Yes	Yes	Yes	Yes	-
Egress Routed v6 ACL on vlans	-	Yes	Yes	Yes	Yes	Yes	-	-
Egress v6 ACL on vlans	-	Yes	Yes	Yes	Yes	Yes	Yes	-
Egress MAC ACL on ports	-	Yes	-	-	Yes	Yes	Yes	-
Egress MAC ACL on vlans	-	Yes	-	-	Yes	Yes	Yes	-
Control Plane ACLs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ingress ADC on ports*	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Bridge-mode SSID

In the AP model using bridge-mode SSIDs, the user VLANs exist on the APs, and wireless traffic is bridged locally to the AP. In the AP with Gateway model using tunnel-mode SSIDs described in the Dynamic Segmentation section, the user VLANs are terminated on the Gateway or the AP, and the wireless traffic is tunneled to the Gateway or bridge locally depending on traffic requirement.

With bridge-mode SSIDs, the AP acts as the policy enforcement point by segmenting the traffic into the appropriate VLAN. The AP is trunked to the access switch, and the VLANs are terminated in the collapsed core. The policy is configured in Central or ClearPass, which provide easy ways to administer and manage policy across an organization from a centralized location. Bridge-mode SSIDs are recommended when customers expect to deploy less than 500 APs and have less then 5000 clients per site.

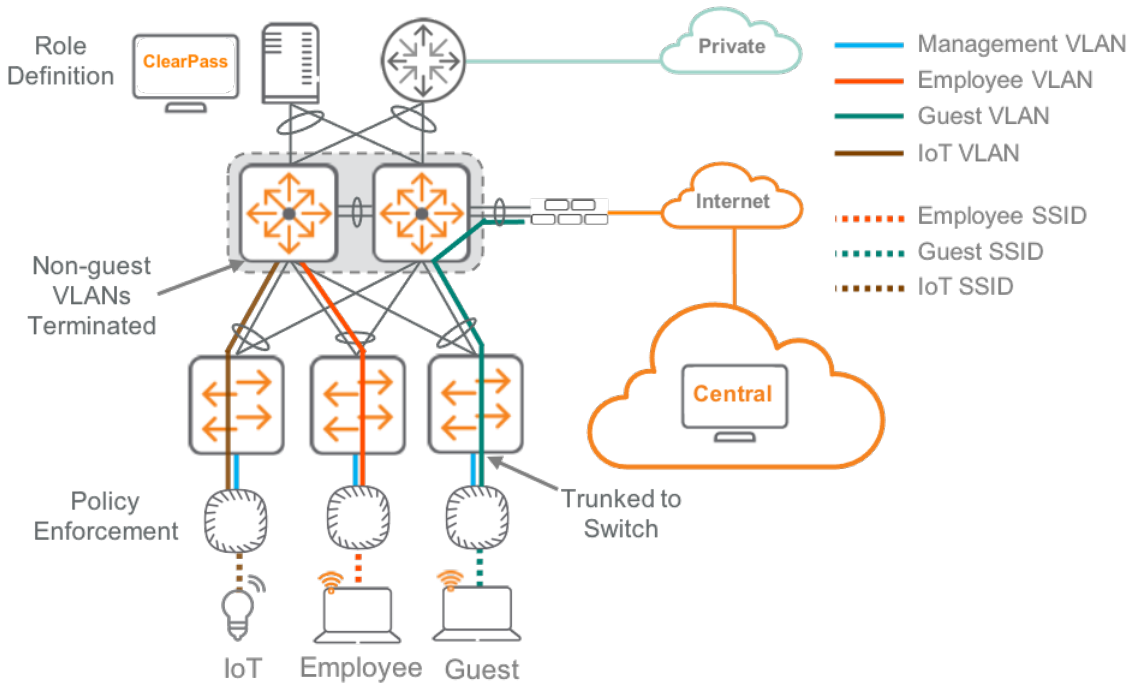

Figure 34: **Bridge-mode in two-tier wired**

Mixed-mode SSID

Aruba also supports a new SSID forwarding mode called mixed-mode. This mode allows a customer to enable both bridge and tunnel forwarding modes on a single SSID without requiring separate SSIDs for each mode. The benefits of a mixed-mode SSID is to reduce the overall number of SSIDs. Fewer SSIDs in a busy network, and the reduced percentage of management/control frames and beacons, help boost the WLAN network performance and simplify overall network management. Mixed-mode SSIDs are recommended when a deployment needs bridged and tunneled traffic and wants to use 802.1X authentication for both tunneled and bridged traffic.

The key thing to note is mixed-mode SSIDs only support 802.1X for authentication. Aruba ClearPass is preferred but not mandatory to deploy mixed-mode SSIDs. Bridged and tunneled VLAN derivation are done using a standard vendor-specific attribute VSA (example filter-id) or an Aruba VSA (example user-role), which is provided by ClearPass or any RADIUS server.

Campus Services Layer

The *Services Layer* is where the operations team interacts with the Connectivity and Policy layers. It provides significant capabilities leveraging AI, ML and location-based services for network visibility and insights into how the network is performing. By leveraging a unified data lake in the cloud, Aruba ESP correlates cross-domain events and displays multiple dimensions of information in context, unlocking powerful capabilities around automated root-cause analysis while providing robust analytics. The primary homes for Services Layer functionality are Central and ClearPass Policy Manager.

Aruba Central

Central is designed to simplify the deployment, management and optimization of WLAN, LAN, VPN and SD-WAN. This capability allows administrators to eliminate the time-consuming and manual process of moving information from management platform to management platform or trying to correlate troubleshooting information across multiple disconnected views. As the single pane of glass for Aruba ESP, the use of integrated AI-based machine learning, IoT device profiling for security and unified infrastructure management accelerates the edge-to-cloud transformation for today's intelligent edge.

Central Key Features

- Cloud-native enterprise campus WLAN software
- AI Insight for WLAN, switching, and SD-WAN
- Advanced IPS/IDS threat defense management
- Mobile application-based network installation
- Unified management for access and WAN edge
- Live Chat and an AI-based search engine
- Cloud, on-premises and as-a-Service options

Central is a cloud-native micro services-based platform that provides the scalability and resiliency needed for mission-critical environments across the distributed edge. Since Central runs in the cloud, it is adaptive, predictable and horizontally scalable with built-in redundancy, unlike an on-premises solution. Central also provides seamless access to ClearPass Device Insights, User eXperience Insights, and Meridian to furnish significant capabilities leveraging AI/ML, and location-based services for network visibility and insights.

ClearPass Policy Manager

ClearPass Policy Manager provides role and device-based secure network access control for IoT, BYOD, corporate devices, as well as employees, contractors and visitors across wired, wireless and VPN infrastructure. With a built-in context-based policy engine, RADIUS, TACACS+, non-RADIUS enforcement using OnConnect, device profiling, posture assessment, onboarding, and visitor access options, ClearPass is unrivaled as a foundation for network security for organizations of any size.

ClearPass also supports secure self-service capabilities, making it easier for end users trying to access the network. Users can securely configure their own devices for enterprise use or Internet access based on admin policy controls. Aruba wireless customers in particular can take advantage of unique integration capabilities such as AirGroup, as well as ClearPass Auto Sign-On (ASO). ASO enables a user's network authentication to pass automatically to their enterprise mobile apps so they can get right to work.

ClearPass Policy Manager Key Features

- Role-based, unified network access enforcement across multi-vendor wireless, wired and VPN networks
- Intuitive policy configuration templates and visibility troubleshooting tools
- Supports multiple authentication/authorization sources (AD, LDAP, SQL)
- Self-service device onboarding with built-in certificate authority (CA) for BYOD
- Visitor access with extensive customization, branding and sponsor-based approvals
- Integration with key UEM solutions for in-depth device assessments
- Comprehensive integration with the Aruba 360 Security Exchange Program

ClearPass is the only policy platform that centrally enforces all aspects of enterprise-grade access security for any industry. Granular policy enforcement is based on a user's role, device type and role, authentication method, UEM attributes, device health, traffic patterns, location, and time-of-day. Deployment scalability supports tens of thousands of devices and authentications which surpasses the capabilities offered by legacy AAA solutions. Options exist for small to large organizations, from centralized to distributed environments.

ClearPass Device Insight

Today's networks have become increasingly more complex, due in part to the rapid adoption of Internet of Things (IoT) devices which are often difficult to detect and manage. In order to leverage the operational efficiencies of mobile and IoT, many organizations are deploying a wide range of devices, without fully understanding the security and compliance implications.

Aruba ClearPass Device Insight provides a full-spectrum of visibility across the network by intelligently discovering and profiling all connected devices. This includes detailed device attributes such as device type, vendor, hardware version, and behavior including applications and resources accessed. This allows organizations to create more granular access policies, reduce security risks and meet key compliance requirements. As a part of Aruba's ClearPass family of industry-leading access control solution, ClearPass Device Insight provides the visibility needed to make better informed network access control decisions. Integration with ClearPass Policy Manager delivers comprehensive policy control and real time enforcement. This makes the visibility provided by ClearPass Device Insight actionable and increases the overall level of security and compliance for all devices connected to the network.

User eXperience Insight

Aruba User Experience Insight (UXI) is a cloud-based service assurance solution that validates network health and troubleshoots problems that affect day-to-day user experience. Ideal for campus and branch environments alike, UXI assumes the role of an end-user, evaluating the performance, connectivity, and responsiveness of network infrastructure as well as internal and external services such as corporate ERM or Office365 applications. This outside-in perspective is presented through a simple, intuitive dashboard that provides a proactive way to solve problems before they impact the business. UXI is easy to configure, deploy and manage, and immediately begins providing insights once sites are online.

Meridian

Aruba Meridian is a cloud-based, software-as-a-service (SaaS) solution that is part of Aruba's location services portfolio. It includes both mobile engagement and asset tracking capabilities. Customers can start with wayfinding, proximity-based notifications or digital asset tracking and add functionality as needed via simple subscription-based licensing. Meridian's cloud-based architecture allows organizations and venues like corporate and university campuses, as well as stadiums, airports, museums, hospitals, and retail stores to easily manage their location services needs from anywhere. The inclusion of proximity-based notifications and analytics makes Meridian the industry's leading full featured location services platform.

Service Capabilities

Some of the key service capabilities of the ESP Campus include Live Upgrade, AI Insights, AI Assist, AirGroup, Air Slice, AirMatch and ClientMatch. The nature of Central as a services platform means capabilities will continue to be added over time without the need for infrastructure upgrades or significant design overhauls for a customer's environment.

Live Upgrade is an Aruba technology that uses telemetry data obtained from the network to understand how a network can be upgraded with the least amount of impact and then, coordinates that upgrade between clients and hardware to minimize the need for maintenance windows and downtime.

AI Insights is a capability in Central specifically built to quickly identify, categorize, and resolve issues that would impact client onboarding, connectivity and network optimization. These insights provide clear descriptions of the detected issue, visualizations of the data, recommended fixes, and contextual data to determine the overall impact. AI Insights uses ML-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 AI Insight cloud instance where network connectivity and performance issues are analyzed by leveraging ML-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 AI Insights to automatically power down when connectivity demand ceases and power up when demand returns. AI Insights uses predictive analytics and ML to identify usage patterns. After a brief learning period, AI Insights can predict when demand stops and when it starts.

AI Assist is the always-on technical assistant which helps augment the network operations team. AI Assist uses event-driven automation to collect and post relevant data for both the internal help desk and the Aruba Technical Assistant Center. Having all the data available about an issue centralized in one place removes the need for network administrators to use multiple analytical tools. Everything about an event is displayed in context, in single views to help resolve problems very quickly.

AirGroup is an Aruba technology that will aide in mDNS and SSDP style discovery protocols across VLANs. AirGroup will also allow for a personal group of these devices that can be access no matter the location or VLAN the client is on. Several technologies can be used together with AirGroup to scale to a customer's needs. AirGroup brings Enterprise controls to technologies not designed for the Enterprise.

Air Slice allows for prioritization of client traffic at the radio level on Aruba's Wi-Fi 6 APs. This technology is transparent to the client, so it has no integration or standards requirements to work with a client unlike older technologies. Air Slice has a tight integration into the DPI firewall capabilities of the AP so Air Slice policies can be created and based off applications instead of ports and IP addresses.

AirMatch provides automated RF optimization by dynamically adapting to the ever-changing RF environment at the network facility. In the ESP solution, the AirMatch service is moved to Central, which is capable of computing and deploying RF allocation to APs across the entire enterprise network. The AirMatch service receives telemetry data from APs for radio measurements, channel range, transmit power range, operational conditions, and local RF events like radar detection or high noise.

ClientMatch is the initial feature that allowed Aruba to be the first networking vendor to offer AI/ML capabilities to their customers. ClientMatch optimizes the client association by continuously scanning the wireless environment and sharing information about the clients and the APs. Based on the dynamic data obtained, clients are steered to the most suitable AP and no software changes are required in the clients to achieve this functionality.

Reference Architectures for Campus

This section includes architectures for small, medium and large networks as well as campuses consisting of multiple different size buildings. The purpose of this section is to show characteristics of a design in order to choose when to use a two-tier versus a three-tier campus architecture. In addition to each scenario, there is a list of features and functions that are recommended for each type of network. It is up to the reader to decide which network scenario best fits their environment. In addition to providing approximate sizing, the detailed recommendation for the set of features will be covered in the respective sections.

Small Campus

The small campus architecture is targeted for organizations supporting up to 5000 users with multiple devices per user. The network could be a single building, a few floors in a larger building, or a group of small buildings located near each other. The wireless network requires a common LAN design which consists of two tiers. The access layer is where wired devices and wireless APs connect to the network. The collapsed core layer acts as a connection point for multiple access-layer switches and where common services are connected to the network.

The following example is a reference design for a small campus consisting of several floors in a building to give the reader general guidelines around sizing requirements. The design includes one main distribution frame (MDF) and combined server room, and several intermediate distribution frames (IDF) across the floors that connects to the MDF using multi-mode fiber. This small campus reference design supports 750 employees and requires 75 APs to provide 2.4 GHz and 5 GHz coverage.

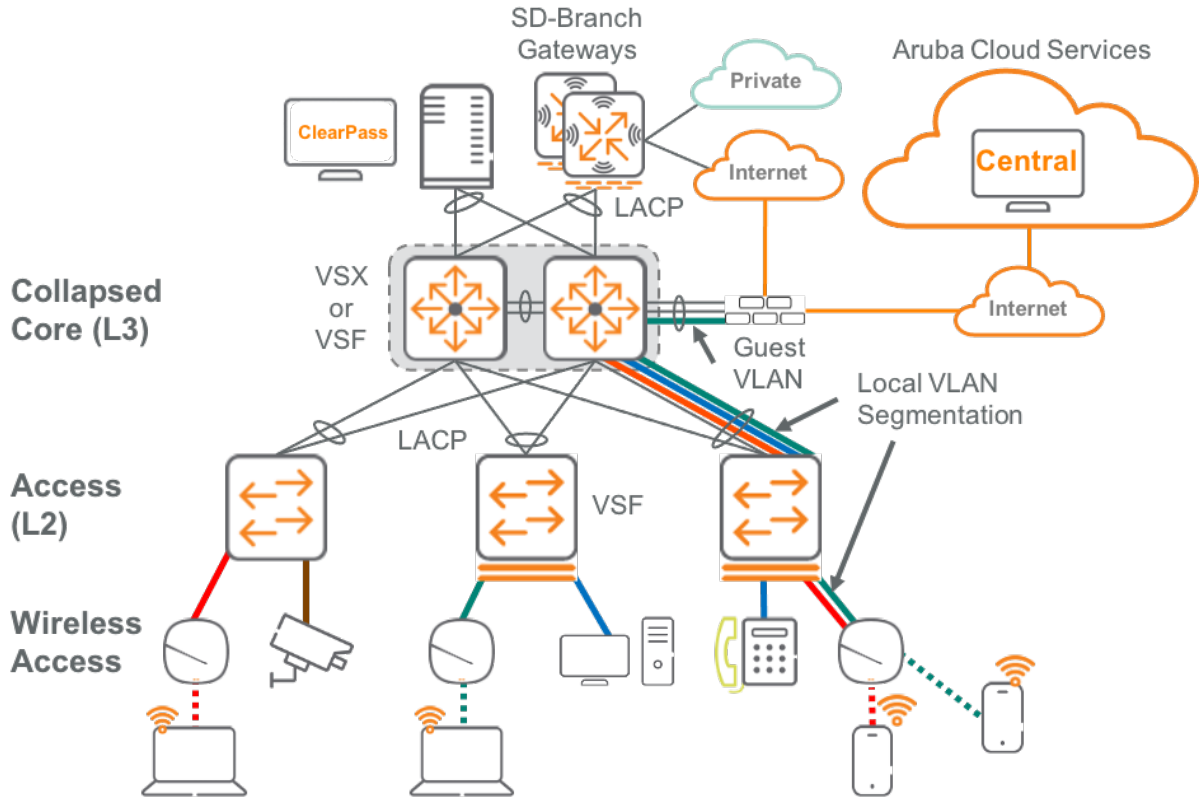

Figure 35: Small campus architecture

Building Characteristics:

- 3 - 4 Floors / 75,000 sq ft Total Size
- 750 x Employees with up to 2500 x Concurrent IPv4 Clients
- 75 x APs
- 1 x Combined Server Room / Wiring Closet (MDF)
- 10 x Wiring Closets (IDF)

This is several floors in a single building with a limited number of access closets and therefore does not require an aggregation layer between the collapsed core and access layer. This building will implement a two-tier network design where the access layer switches and services connect directly to a collapsed core / aggregation layer. This two-tier network design can also accommodate small buildings with a larger sq. footage and additional floors if required.

LAN Components

The small building in this scenario includes various LAN components which are deployed in the user space, access closet or aggregation closet. The tables below summarize the roles and provides general guidance on the number of devices recommended.

Table 6: Switch role recommendations

switch Roles	Description	Notes
Core	Building Core Switch	Not required
Collapsed Core	Aggregation/Collapsed Core switch SFP/SFP+/SFP28 (Access layer interconnects)	2 x Required (Collapsed Core)
Access Layer	Access Switch SFP/SFP+ (Core / aggregation layer interconnects) 10/100/1000BASE-T/HPE SmartRate (Edge ports)	Minimum of 2 per IDF recommended
MDF	Provides connectivity for compute resources	Optional

To accommodate the AP and client counts for this scenario, a Gateway is not required in this design.

Table 7: Wireless role recommendations

Wireless Roles	Description	Notes
Access Points	Wi-Fi 6 or Wi-Fi 5 Access Points	15 x Required

NOTE:

To accommodate the AP and client counts for this scenario, a Gateway is not required in this design.

Local Area Network

This small campus scenario has a two-tier network with a collapsed core design. The collapsed core switches will connect to the access switches with Layer 2 MC-LAG links and provide all routing between VLANs. The collapsed core switch can provide connectivity to an optional switch stack for local compute resources. The size of the small building may not warrant in-building local compute or a dedicated MDF switch.

The recommended two-tier design calls for switch redundancy for both collapsed core and aggregation access. This can be achieved using Virtual switch Extension (VSX) on the collapsed core devices and LACP from the access up to the VSX pair. The recommended access design would be to use Virtual Stacking Framework (VSF) between two or more switches per IDF. Each access switch should have a single LACP uplink to the upstream collapsed core switch allowing traffic to flow through either member if one goes down. The access switch stacks would provide power for access points and other

PoE devices. APs should be alternated between stack members so wireless clients will have coverage if one switch is off-line, the second switch can still provide power/connectivity to AP's. This is done by distributing AP connections between switches so that if one switch is down AP coverage will not affect clients.

The optional MDF switch has similar redundancy considerations as the core/aggregation switch stack. While PoE is not needed in the computer room, consider the impact of a computer room switch outage. In most cases, the cost of an outage is sufficient that having redundant computer room switches is highly desirable. This is especially true if the devices which will be connected to the switches have the ability to be dual-attached to minimize the impact of a switch failure.

With a two-tier network it can be difficult to identify which switch should be used where and why. In a two-tier network there are two layers to focus on the collapsed core and the access. The ideal switches for a collapsed core would be the 8360 due to port count and feature set. It is important to note that low density 8360's only have two 40G uplink ports so they should be used as VSX interconnects and one of the downlink ports should be used as a connection to other services such as firewalls and servers.

Table 8: Collapsed Core switch models, benefits and considerations

Collapsed Core	Benefits	Considerations
8400	Modular good for growth Scale (786K MAC / 1M routes)	High scale numbers, high cost
6400	Modular good for growth Scale (32K MAC / 64K routes)	Low scale
8360 (Recommended)	High scale, VXLAN capable Scale (212K MAC / 600K routes)	Low port count options (16/24 port options)
8325	High port count, VXLAN capable Scale (98K MAC / 131K routes)	No low port count options
8320 (Alternative)	High port count Scale (98K MAC / 131K routes)	No low port count options, not VXLAN capable

Table 9: Access switch models, benefits and considerations

Access	Benefits	Considerations
6400	Modularity, management redundancy, VNBT/UBT capable	High cost for small deployment
6300M (Recommended)	Redundant power supplies, VNBT/UBT capable	Fans can be loud for a front desk

Access	Benefits	Considerations
6300F (Alternative)	Very quiet, VNBT/UBT capable	No redundant power supplies
6200 (Alternative)	Very quiet, UBT capable	No redundant power supplies No route only ports
6100 (Alternative)	Low port count / quiet / portable	No advanced feature set Not UBT/VNBT capable

NOTE:

Alternate Switches are available depending on the requirements of the environment, check the data sheet or consult a local SE or partner for more details.

Table 10: Wired features in a two-tier network

Wired Features	Collapsed Core Layer	Access Layer
IP Design Features / Functions		
IP Routing	Yes	Optional
PIM BSR	Optional	No
PIM cRP	Optional	No
PIM DR	Optional	No
MSDP	Optional	No
IGMP	No	Optional
Base Configuration		
NTP/Time zones	Yes	Yes
sFlow	Yes	Recommended
NAE	Yes	Yes
Device hardening	Yes	Yes
Traffic Management		
QoS	Yes	Yes
Spanning Tree	Yes	Yes
Loop protect	Yes	Yes
Network Resiliency		
VSX	Yes	No
VSF	No	Recommended

Wired Features	Collapsed Core Layer	Access Layer
Access Security		
Radius	No	Recommended
User-Based Tunneling	No	Optional
Colorless ports	No	Recommended
Virtual Network-Based Tunneling	Optional	Optional
VRF	Optional	Optional

Table 11: Wireless features for APs

Wireless Features	Access Points
IP Design Features / Functions	
DMO	Optional
Base Configuration	
NTP/Time zones	Yes
AP Firewall	Optional
Device hardening	Yes
Traffic Management	
QoS	Optional
AirMatch	Recommended
ClientMatch	Recommended
Air Slice	Optional
Spanning Tree	Yes
Access Security	
Radius	Recommended

The recommended features depend on certain applications to enable their functionality. ClearPass, NetEdit and Central are three applications within the ESP architecture. The small campus requires ClearPass to enable device/client authentication and device visibility. In addition to authentication, ClearPass is required for user-based tunneling and colorless ports. It can also be used for VNBT but is not required. Central is the management platform used to monitor, edit and configure devices within the network and is highly recommended. However, there are cases where cloud access is not an available, so NetEdit is optional for wired management.

Table 12: Services recommendations

Services	Description	Notes
Aruba ClearPass	Virtual Appliance	Recommended (On-premises)
Aruba NetEdit	Virtual Appliance	Optional for wired (On-premises)
Aruba Central	Cloud Appliance	Recommended (Cloud)

Medium Campus

The medium campus architecture is targeted for organizations supporting 5000 to 15000 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 LAN design which consists of three tiers. The access layer is where wired devices and wireless APs connect to the network. The access aggregation layer acts as a connection point for multiple access-layer switches. The core / services layer interconnects aggregation-layer switches from multiple buildings or multiple floors in a building and it is where the common services are connected to the network. The three-tier design is used when the number of aggregation switch pairs exceeds two or the layout of the physical wiring plant makes more sense to connect everything to a central core.

The following example is a reference design for a medium campus consisting of ten to twelve floors in a building to give the reader general guidelines around sizing requirements. The building includes a data center which connects to the core/services switch in the main distribution frame (MDF). Each floor includes two or more intermediate distribution frames (IDFs) which connect to the MDF. This medium campus reference design can support up to 7,000 employees and requires up to 700 APs to provide full 2.4 GHz and 5 GHz coverage.

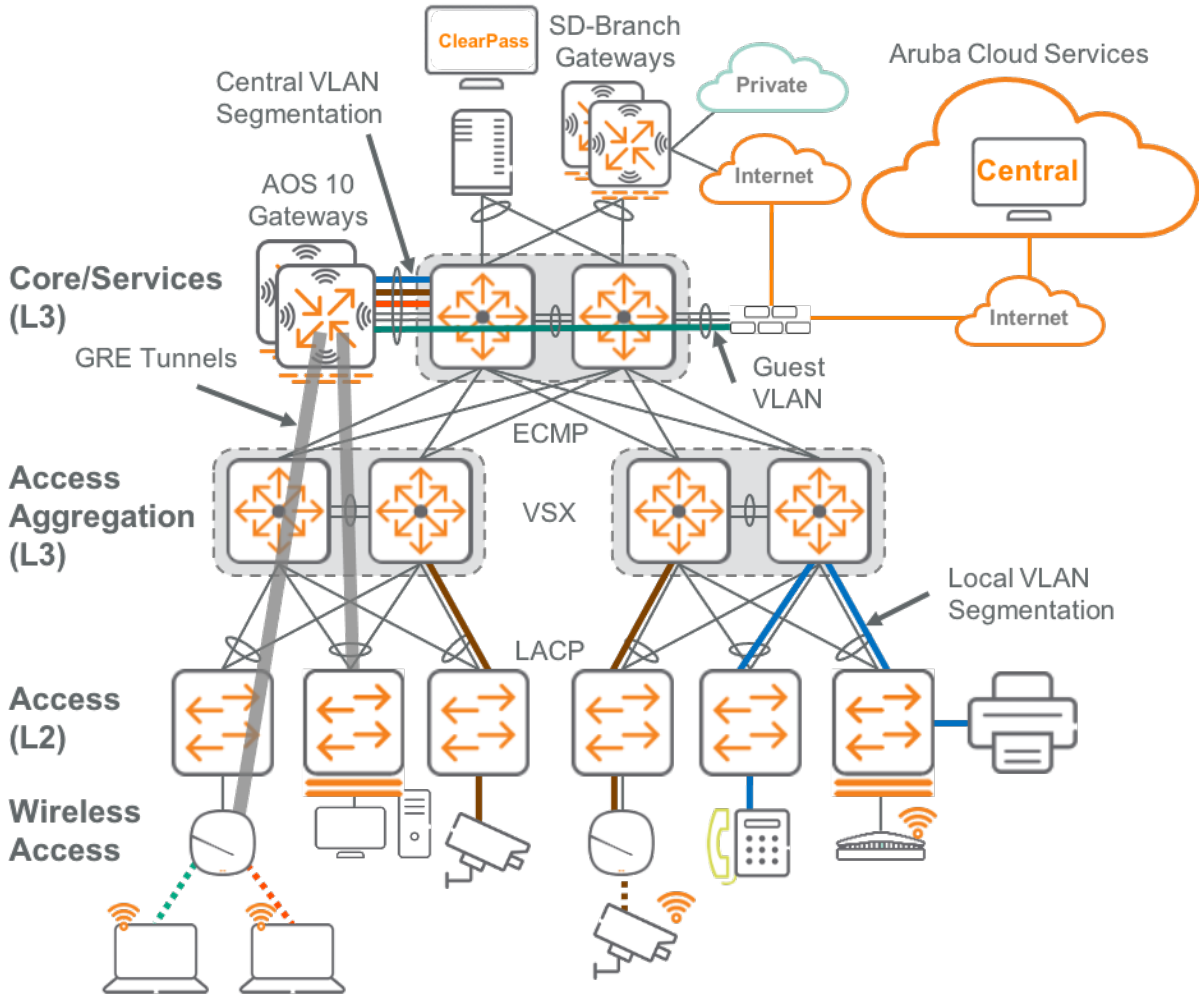

Figure 36: Medium campus architecture

Building Characteristics:

- 15 - 20 Floors / 700,000 sq ft Total Size
- 7,000 x Employees with up to 25,000 x Concurrent IPv4 Clients
- 700 x APs
- 2 x Gateways
- 1 x Computer Room
- 1 x MDF
- 2 or 3 x IDFs per floor/building

As this building implements a structured wiring design using MDFs and IDFs, an aggregation layer to connect the access layer is required. This building will implement a three-tier modular network design where the access layer switches connect via aggregation layer switches in each MDF that connect directly to the core. This modular network design also includes a core / services layer with enough capacity for the computer room servers.

LAN Components

The medium building in this scenario includes various LAN components which are deployed in the user space, access closet, aggregation closet or building core location. The tables below summarize the roles and provides general guidance on the number of devices recommended.

Table 13: Switch role recommendations

Switch Roles	Description	Notes
Core / Services	Building Core / Services Switch SFP/SFP+/QSFP+ (Aggregation Layer Interconnects) SFP/SFP+ (Module Connectivity)	2 x Required
Aggregation	Aggregation Switch SFP/SFP+/SFP28/QSFP+ (Core and Access Layer Interconnects)	2 x Per Floor
Access	Access Layer Switch SFP/SFP+/SFP28 (Aggregation Layer Interconnects) 10/100/1000BASE-T/HPE SmartRate (Edge ports)	Minimum of 2 per IDF recommended
MDF	Provides connectivity for compute resources	Optional

The wireless components are either deployed in the aggregation closet or the server room. To accommodate the AP and client counts for this scenario, a single cluster of gateways is required. The number of cluster members is determined by the gateways selected (see platform suggestions). For redundancy, the gateway cluster consists of a minimum of two gateways with each member providing adequate capacity and performance to operate the wireless network in the event of a single gateway failure. For this scenario, the Gateways are deployed within a computer room and connect directly to the core / services switches.

Table 14: Wireless role recommendations

Wireless Roles	Description	Notes
Gateways	Hardware appliances	2 x Minimum (Clustered)
Access Points	Wi-Fi 6 or Wi-Fi 5 APs	700 x Required

Local Area Network

This scenario has a three-tier network providing dedicated access, aggregation and core layers. The core layer will consist of a pair of devices to provide redundancy and eliminate single points of failure. The aggregation layer consists of two pair of switches each providing redundant connectivity for connections to both the core and access layers. The aggregation switches will connect to the access switches with Layer 2 links and will provide any and all routing between VLANS. The aggregation layer switches will connect to the core switches via Layer 3 links. The core switch can provide connectivity to an optional switch stack for any local compute resources. The size of the medium building may not warrant having any in-building local compute or a dedicated computer room switch.

The recommended core/aggregation design calls for switch redundancy which can be achieved using LACP from the access to a Virtual switch Extension (VSX) Aggregation. The recommended access switch design would be to use one or more switches per IDF in a stacking configuration. The switch stacks would need to provide enough power for access points and other PoE devices as well as provide enough Ethernet interfaces for wired systems. Stacking is recommended to build fault tolerant designs so that if one switch is off-line, the is still connectivity to access points and the building core/aggregation switches. Each access switch should have distributed AP connections between switches so that if one switch is down AP coverage will not affect clients. Uplink connectivity from the access to the aggregation layer would be provided by uplink ports (using ports from different switches in the stack) configured in an LACP group.

The optional services switch has similar redundancy considerations as the aggregation switch stack. While PoE is not needed the impact of a computer room switch outage needs to be considered. In most cases, the cost of an outage is sufficient that having redundant computer room switches is highly desirable. This is especially true if the devices which will be connected to the switches have the ability to be dual-attached, we can minimize the impact of a switch failure.

The three-tier network has many available switch options. The core layer of the network should take into account the device feature set and scale. There are additional factors like modularity that come into play when expanding the network over time.

Table 15: Core switch models, benefits and considerations

Core	Benefits	Considerations
8400	Redundant power / Modular expansion Scale (1M Routes)	Good for meeting a modularity requirement Good for planning for future expansions
6400 (Alternative)	Redundant power / Modular expansionScale (64K routes)	Good for meeting a modularity requirement for a site that doesn't plan on expanding
8360 (Alternative)	Redundant powerScale (600K routes)	Only 1 RU so there is a set number of buildings/services that can be supported

Core	Benefits	Considerations
8325 (Recommended)	Redundant powerScale (131K routes)	Only 1 RU so there is a set number of buildings/services that can be supported
8320	Redundant powerScale (131K routes)	Only 1 RU so there is a set number of buildings/services that can be supported

Aggregation layer switches are used to consolidate access switches to a common set of devices the main factors in deciding which aggregation switch to go with are feature set, scale and port count.

Table 16: Aggregation switch models, benefits and considerations

Aggregation	Benefits	Considerations
8400	Redundant power / Modular expansion Scale (786K MAC / 1M routes)	Potentially high port count for limited number of switches and very high scale
6400	Redundant power / Modular expansion Scale (32K MAC / 64K routes)	Potentially high port count for limited number of switches
8360 (Alternative)	Redundant power, VXLAN ready Scale (212K MAC / 600K routes)	Ensure low port count options are not selected
8325 (Recommended)	Redundant power, VXLAN ready Scale (98K MAC / 131K routes)	
8320 (Alternative)	Redundant power Scale (98K MAC / 131K routes)	No VXLAN

Table 17: Access switch models, benefits and considerations

Access	Benefits	Considerations
6400	Modularity, management redundancy	High cost for medium deployment
6300M (Recommended)	Redundant power supplies, VNBT/UBT capable	Fans can be loud for a front desk
6300F (Alternative)	Very quiet, VNBT/UBT capable	No redundant power supplies
6200 (Alternative)	Very quiet, VNBT/UBT capable	No redundant power supplies No route only ports Not VXLAN Ready

Access	Benefits	Considerations
6100 (Alternative)	Low port count, quiet, portable	No advanced feature set Not VNBT/UBT capable

NOTE:

Alternate Switches are available depending on the requirements of the environment, check the data sheet or consult a local SE or partner for more details.

Table 18: Wired features in a three-tier network

Wired Features	Aggregation Layer	Access Layer
IP Design Features / Functions		
IP Routing	Yes	Optional
PIM BSR	Optional	No
PIM cRP	Optional	No
PIM DR	Optional	No
MSDP	Optional	No
IGMP	No	Optional
Base Configuration		
NTP/Time zones	Yes	Yes
sFlow	Yes	Recommended
NAE	Yes	Yes
Device hardening	Yes	Yes
Traffic Management		
QoS	Yes	Yes
Spanning tree	Yes	Yes
Loop protect	Yes	Yes
Network Resiliency		
VSX	Yes	No
VSF	No	Recommended
Access Security		
Radius	No	Yes
User-Based tunneling	No	Recommended

Wired Features	Aggregation Layer	Access Layer
Colorless ports	No	Recommended
Virtual Network-Based Tunneling	No	Optional
VRF	Optional	Optional

Table 19: Wireless features for Gateways

Wireless Features	Gateway
IP Design Features / Functions	
DMO	Optional
Base Configuration	
NTP/Time zones	Yes
Gateway firewall	Optional
Device hardening	Yes
Traffic Management	
QoS	Recommended
AirMatch	Recommended
ClientMatch	Recommended
Air Slice	Recommended
Spanning Tree	Yes
Network Resiliency	
Clustering	Recommended
Access Security	
Radius	Recommended
User-Based Tunneling	Optional

The recommended features depend on certain applications to enable their functionality. ClearPass, NetEdit and Central are three applications within the ESP architecture. The small campus requires ClearPass to enable device/client authentication and device visibility. In addition to authentication, ClearPass is required for user-based tunneling and colorless ports. It can also be used for VNBT but is not required. Central is the management platform used to monitor, edit and configure devices within the network and is highly recommended. However, there are cases where cloud access is not an available, so NetEdit is optional for wired management.

Table 20: Services recommendations

Services	Description	Notes
Aruba ClearPass	Virtual appliance	Required (On-premises)
Aruba NetEdit	Virtual appliance	Optional for wired (On-premises)
Aruba Central	Cloud appliance	Recommended (Cloud)

Large Campus

The large campus architecture is targeted for organizations supporting more than 15000 users with multiple devices per user. The network could be a single building, or a group of large buildings located near each other. The wireless network requires a common LAN design which consists of three tiers. The access layer is where wired devices and wireless APs connect to the network. The access aggregation layer acts as a connection point for multiple access-layer switches. The standalone core layer interconnects aggregation-layer switches from multiple buildings or multiple floors in a building. The dedicated services aggregation layer is where the common services are connected to the network.

The following example is a reference design for a large campus consisting of 20 or more floors in a building to give the reader general guidelines around sizing requirements. The building includes a data center which connects via single-mode fiber to a main distribution frame (MDF) on each floor. Each floor includes three intermediate distribution frames (IDFs) which connect to the MDF via multi-mode fiber. The large campus supports up to 20,000 employees and requires up to 2000 x APs to provide full 2.4 GHz and 5 GHz coverage.

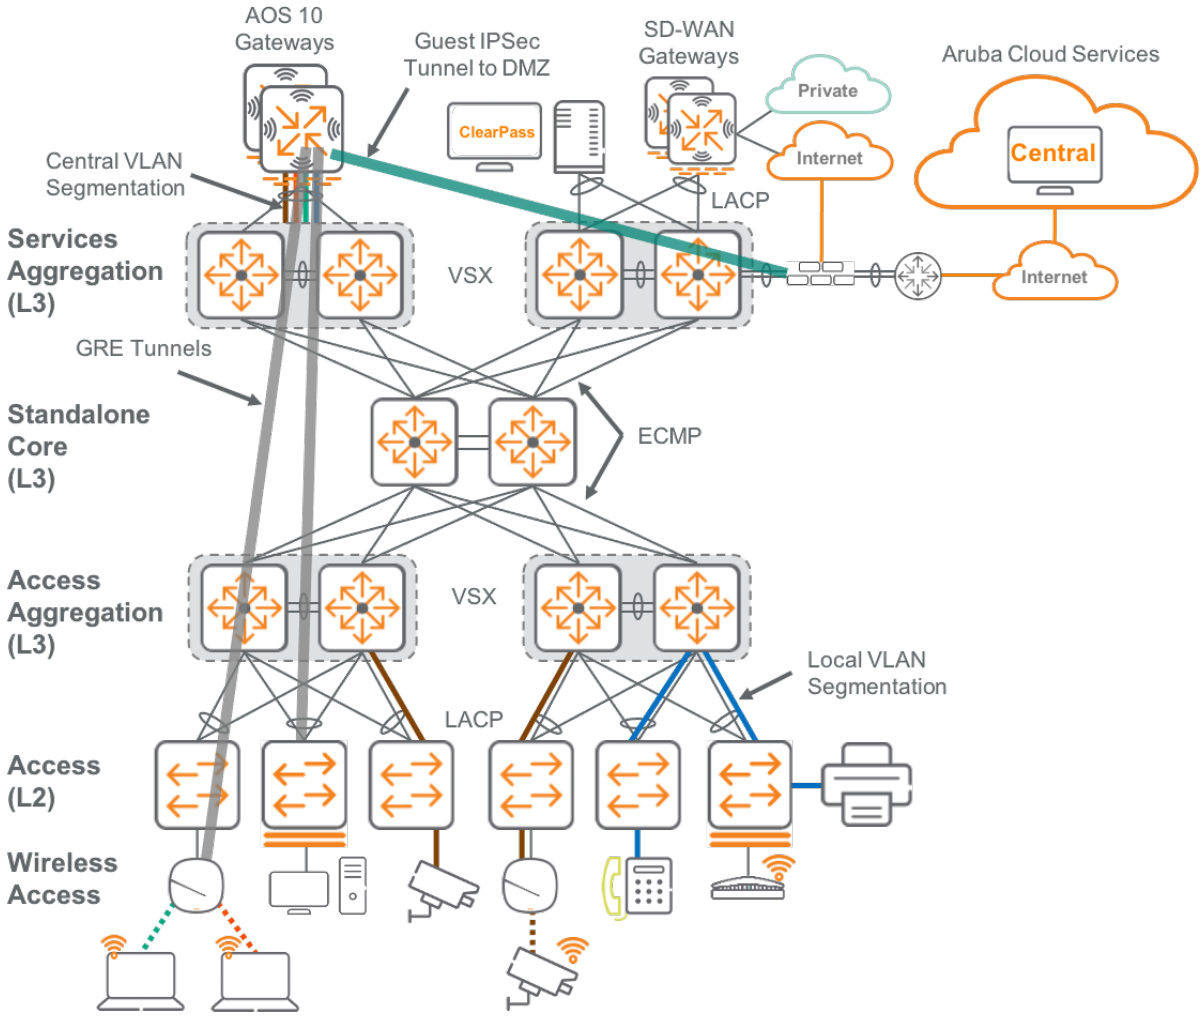

Figure 37: Large campus architecture

Building Characteristics:

- 30+ Floors / 1,800,000 sq ft total size
- 18,000 x Employees with up to 65,000 x Concurrent IPv4 clients
- 1800 x APs
- 4 x Gateways
- 1 x Computer room
- 1 x MDF
- 3 x IDFs per floor

As this building implements a structured wiring design using MDFs and IDFs, an aggregation layer to connect the access layer is required. This building will implement a three-tier modular network design where the access layer switches connect via aggregation layer switches in each MDF that connect directly to the core. For scaling, aggregation and fault domain isolation, this modular network design also includes an additional service aggregation layers for the MDF for wireless modules.

LAN Components

The large building in this scenario includes various LAN components which are deployed in the user space, access closet, aggregation closet or building core location. The tables below summarize the roles and provides general guidance on the number of devices recommended.

Table 21: Switch role recommendations

Switch Roles	Description	Notes
Core	Building Core Switch SFP/SFP+/QSFP+ (Aggregation layer interconnects) SFP/SFP+ (Module connectivity)	2 x Required
Aggregation	Aggregation Switch SFP/SFP+/SFP28/QSFP+ (Core and access layer interconnects)	2 x Per floor
Access	Access Switch SFP/SFP+/SFP28 (Aggregation layer interconnects) 10/100/1000BASE-T/HPE SmartRate (Edge ports)	Minimum of 2 per IDF recommended
MDF	Provides connectivity for compute resources	Optional

This scenario includes various wireless components which are either deployed in the aggregation layer or server room. To accommodate the AP and client counts for this scenario, a single cluster of gateways is required. The number of cluster members determined by the gateway that is selected (see platform suggestions). For redundancy, the gateway cluster consists of a minimum of two gateways with each member providing adequate capacity and performance to operate the wireless network in the event of a single Gateway failure. For this scenario, the Gateways are deployed within a computer room and connect directly to the core or computer room aggregation switches.

Table 22: Wireless role recommendations

Wireless Roles	Description	Notes
Gateways	Hardware appliances	4 x Minimum (Clustered)
Aruba Access Points	Wi-Fi 6 or Wi-Fi 5 APs	1800 x Required

Local Area Network

This scenario has a three-tier network providing dedicated access, aggregation, and core layers. The wireless network also has a dedicated service block providing connectivity for the Gateway cluster. The aggregation layer consists of two pair of switches with each pair providing redundant connectivity for connections to both the core and access layers. The core layer will consist of a pair of devices to provide

redundancy and eliminate single points of failure. The aggregation switches will connect to the access switches with layer-two links and will provide any and all routing between VLANS. The aggregation layer switches will connect to the core switches via layer 3 links. The core switch will also provide connectivity to other service blocks which likely includes connectivity to a data-center, internet edge, and WAN edge service blocks.

With a core layer that is entirely layer 3 connected, the network can leverage equal cost multipath (ECMP) routing to provide for connectivity between core devices as well as to aggregation and other service blocks. The aggregation layer will provide high-availability using VSX. VSX will allow for the elimination of spanning-tree at the aggregation layer and will allow for active/active forwarding to/from access-layer devices.

The recommended access switch design would be to use one or more switches per IDF in a stacking configuration. The switch stacks would need to provide enough power for access points and other PoE devices as well as provide enough Ethernet interfaces for wired systems. Stacking is recommended to build fault tolerant designs so that if one switch is off-line, the is still connectivity to access points and the building core/aggregation switches. Each access switch should have distributed AP connections between switches so that if one switch is down AP coverage will not affect clients. Uplink connectivity from the access to the aggregation layer would be provided by uplink ports (using ports from different switches in the stack) configured in an LACP group.

Large facilities may have an on-site data center. Data Center design is beyond the scope of this document. However, connectivity from the Campus network to the data-center will be include in this design. Fundamentally, connectivity to the data-center has very similar requirements to connectivity to other service blocks including the WAN edge or the internet edge.

NOTE:

Please refer to the ESP Data Center Design Guide for details on Aruba's Data Center designs.

The three-tier network has many available switch options. The core layer of the network should take into account the device feature set and scale. There are additional factors like modularity that come into play when expanding the network over time.

Table 23: Core switch models, benefits and considerations

Core	Benefits	Considerations
8400 (Recommended)	Redundant Power / Modular expansion Scale (1M routes)	
6400	Redundant Power / Modular expansion Scale (64K routes)	Will not scale in large environments
8360 (Alternative)	Redundant power Scale (600K routes)	Low port count

Core	Benefits	Considerations
8325	Redundant power Scale (131K routes)	Low port count / Lower scale compared to recommended alternatives
8320	Redundant power Scale (131K routes)	Low port count / Lower scale compared to recommended alternatives

Aggregation layer switches are used to consolidate access switches to a common set of devices the main factors in deciding which aggregation switch to go with are feature set, scale and port count.

Table 24: Aggregation switch models, benefits and considerations

Aggregation	Benefits	Considerations
8400	Redundant power / Modular expansion Scale (786K MAC / 1M routes)	Potentially high port count for limited number of switches and very high scale
6400	Redundant power / Modular expansion Scale (32K MAC / 64K routes)	Potentially high port count for limited number of switches
8360 (Recommended)	Redundant power, VXLAN capable Scale (212K MAC / 600K routes)	High scale is good for UBT enabled environments as a service aggregation switch Can be used as an Access aggregation, but the scale is high
8325 (Recommended)	Redundant power, VXLAN capable Scale (98K MAC / 131K routes)	Good for Access aggregation, but would be low on scale for Service aggregation
8320 (Alternative)	Redundant power Scale (98K MAC / 131K routes)	This would only be an alternative for the Access aggregation and is not VXLAN capable

Table 25: Access switch models, benefits and considerations

Access	Benefits	Considerations
6400	Modularity, management redundancy Good for high density environments	Port count per floor
6300M (Recommended)	Redundant power supplies, VNBT/UBT capable	Fans can be loud for an open office
6300F (Alternative)	Very quiet, VNBT/UBT capable	No redundant power supplies

Access	Benefits	Considerations
6200 (Alternative)	Very quiet, VNBT/UBT capable	No redundant power supplies No route only ports Not VXLAN capable
6100 (Alternative)	Low port count, quiet, portable	No advanced feature set Not VNBT/UBT capable

NOTE:

Alternate Switches are available depending on the requirements of the environment, check the data sheet or consult a local SE or partner for more details.

Table 26: Wired features in a three-tier network

Wired Features	Collapsed Core Layer	Access Layer
IP Design Features / Functions		
IP Routing	Yes	Optional
PIM BSR	Optional	No
PIM cRP	Optional	No
PIM DR	Optional	No
MSDP	Optional	No
IGMP	No	Optional
Base Configuration		
NTP/Time zones	Yes	Yes
sFlow	Yes	Recommended
NAE	Yes	Yes
Device hardening	Yes	Yes
Traffic Management		
QoS	Yes	Yes
Spanning Tree	Yes	Yes
Loop protect	Yes	Yes
Network Resiliency		
VSX	Yes	No
VSF	No	Recommended
Access Security		

Wired Features	Collapsed Core Layer	Access Layer
Radius	No	Yes
User Based Tunneling	No	Recommended
Colorless ports	No	Recommended
Virtual Network-Based Tunneling	No	Optional
VRF	Optional	Optional

Table 27: Wireless features for Gateways

Wireless Features	Gateway
IP Design Features / Functions	
DMO	Optional
Base Configuration	
NTP/Time zones	Yes
Gateway firewall	Optional
Device hardening	Yes
Traffic Management	
QoS	Recommended
AirMatch	Recommended
ClientMatch	Recommended
Air Slice	Recommended
Spanning Tree	Yes
Network Resiliency	
Clustering	Recommended
Access Security	
Radius	Recommended
User-Based Tunneling	Optional

The recommended features depend on certain applications to enable their functionality. ClearPass, NetEdit and Central are three applications within the ESP architecture. The small campus requires ClearPass to enable device/client authentication and device visibility. In addition to authentication, ClearPass is required for user-based tunneling and colorless ports. It can also be used for VNBT but is not required. Central is the management platform used to monitor, edit and configure devices within the network and is highly recommended. However, there are cases where cloud access is not an available, so NetEdit is optional for wired management.

Table 28: Services recommendations

Services	Description	Notes
Aruba ClearPass	Virtual Appliance / Hardware appliance	Required (On-premises)
Aruba NetEdit	Virtual Appliance	Optional for wired (On-premises)
Aruba Central	Cloud Appliance	Recommended (Cloud)

Capacity Planning

The following section provides switching and Gateway capacity planning guidance for the ESP Campus reference architectures. The architectures were thoroughly tested in an end-to-end solution environment that incorporates best practices deployment recommendations, applications and load profiles that represent production environments.

The following tables provides validated values for capacity planning of the ESP Campus design.

Core and Aggregation switch Scaling

Please refer to the Aruba product data sheets for detailed specifications not included in this guide.

Aruba Campus Core and Aggregation switches

Table 29: 8XXX Core and Aggregation switch

Feature	84xx	8360	8325	8320
#VLANs	4,094	4,094	4,040	4,040
#ACLs	16,000 with 64,000 entries per ACL	4,000 with 8,000 entries per ACL	512 with 2,304 entries per ACL	4,000 with 14,336 entries per ACL
ACL Entries ingress	IPv4: 512,000 IPv6: 98,304 MAC: 98,304	IPv4: 65,536 IPv6: 16,384 MAC: 65,536	IPv4: 2,304 IPv6: 2,304 MAC: 2,304	IPv4: 14,336 IPv6: 7,168
ACL Entries egress	IPv4: 198,656	IPv4: 8,192 IPv6: 2,048 MAC: 8,192	IPv4: 2,304 IPv6: 256	IPv4: 256 IPv6: 255

Feature	84xx	8360	8325	8320
MAC	768,000	212,992	98,304	98,304
ARP	IPv4: 756,000 IPv6: 524,000	IPv4:145,780 IPv6:145,780	IPv4: 120,000 IPv6: 52,000	IPv4:120,000 IPv6:52,000
Routing	IPv4: 1,011,712 IPv6: 524,288 v4+v6: 1,011,712	IPv4: 606,977 IPv6: 630,784 v4+v6: 606,977	IPv4: 131,072 IPv6: 32,732 v4+v6: 163,796	IPv4: 131,072 IPv6: 32,732 v4+v6: 163,796
IGMP/MLD	I: 32,767 M: 32,767	I: 7,000 M: 7,000	I: 4,094 M: 4,094	I: 4,094 M: 4,094
Multicast routes	IPv4: 32,767 IPv6: 32,767	IPv4: 7,000 IPv6: 7,000	IPv4: 4,094 IPv6: 4,094	IPv4: 4,094 IPv6: 4,094
Active GW	IPv4: 4,094 IPv6: 4,094 v4+v6: 4,094	IPv4: 1,024 IPv6: 1,024 v4+v6: 1,026	IPv4: 4,040 IPv6: 4,040 v4+v6: 4,040	IPv4: 4,040 IPv6: 4,040 v4+v6: 4,040
#LAGs	256 16p per LAG	52 16p per LAG	56 (32 for JL627A) 16p per LAG	54 (32 for JL759A) 16p per LAG
#VRFs	256	256	256	256

Table 30: 6XXX Aggregation switch

Feature	6400	6300
#VLANs	4,094	4,094
#ACLs	4,000 with8,000 entries per ACL	4,000 with 8,000 entries per ACL
ACL Entries ingress	IPv4: 64,000 IPv6: 64,000 MAC: 64,000	IPv4: 20,480 IPv6: 5,120 MAC: 20,480
ACL Entries egress	IPv4: 64,000 IPv6: 20,460 MAC: 64,000	IPv4: 8,192 IPv6: 2,048 MAC: 8,192
MAC	32,768	32,768
ARP	IPv4: 49,152 IPv6: 49,152	IPv4: 49,152 IPv6: 49,152
Routing	IPv4: 61,000 IPv6: 61,000 v4+v6: 65,536	IPv4: 61,000 IPv6: 61,000 v4+v6: 65,536
IGMP/MLD	I: 7,000 M: 7,000	I: 8,192 M: 8,192
Multicast routes	IPv4: 8,192 IPv6: 8,192	IPv4: 8,192 IPv6: 8,192
Active GW	IPv4: 1,024 IPv6: 1,024 v4+v6: 1,024	IPv4: 1,024 IPv6: 1,024 v4+v6: 1,024
#LAGs	256 16p per LAG	52 16p per LAG
#VRFs	256	256

Access switch Scaling

Please refer to the Aruba product data sheets for detailed specifications not included in this guide.

Aruba Campus Access switches

Table 31: 6XXX Access switch

Switch	VLANs	ACLs	ACL Entries ingress	ACL Entries egress	MAC Table	UBT Clients per port	UBT Clients per system
6400	4,094	4000 with 8000 entries per ACL	IPv4: 64,000 IPv6: 64,000 MAC: 64,000	IPv4: 64,000 IPv6: 20,460 MAC: 64,000	32,768	256	1,024
6300	4,094	4000 with 8000 entries per ACL	IPv4: 20,480 IPv6: 5120 MAC: 20,480	IPv4: 8,192 IPv6: 2,048 MAC: 8,192	32,768	256	1,024
6200	2,048	4000 with 8000 entries per ACL	IPv4: 5,120 IPv6: 1,280 MAC: 5,120	IPv4: 2,048 IPv6: 512 MAC: 2,048	16,000	128	1,024

Gateway Scaling

Please refer to the Aruba product data sheets for detailed specifications not included in this guide.

Aruba Gateways

Aruba Indoor Access Points

Table 32: 7XXX Gateway

Gateway	UBT Tunnels	APs	Clients	VLANs	MACs	IPsec	Active FW sessions	ACLs	Roles
7280	34,816	2,048	32,000	4,094	131,072	32,768	1,940,717	2,680	1,318
7240	34,816	2,048	32,000	4,094	131,047	32,768	1,917,569	2,680	1,318
7240XM	34,816	2,048	32,000	4,094	131,047	32,768	1,917,569	2,680	1,318
7220	17,408	1,024	32,000	4,094	130,899	24,576	1,896,839	2,680	1,318
7210	8,704	512	16,000	4,094	130,363	16,384	1,834,959	2,680	1,318
7205	4,352	256	8,000	4,094	130,764	8,192	1,910,147	2,680	1,318
7030	1,088	64	4,000	4,094	65,536	4,096	63,783	2,680	1,318
7024	544	32	2,000	4,094	65,531	2,048	64,171	2,680	1,318
7010	544	32	2,000	4,094	65,529	2,048	63,840	2,680	1,318
7008	272	16	1,000	4,094	65,530	1,024	63,714	2,680	1,318
7005	272	16	1,000	4,094	65,536	1,024	64,844	2,680	1,318

Table 33: 9XXX Gateway

Gateway	UBT Tunnels	APs	Clients	VLANs	MACs	IPsec	Active FW sessions	ACLs	Roles
9004	544	32	2,000	4,094	16,384	2,048	64,000	2,680	1,316
9012	544	64	4,000	4,094	16,384	2,048	64,000	2,680	1,316

Summary

The flow of information is a critical component to a well-run organization. The ESP Campus provides a prescriptive solution based on best practices and validated topologies to keep information moving seamlessly and securely. This allows an organization to build a robust network that easily accommodates their technical requirements without sacrificing capacity.

Whether users are located at a large campus location or at a smaller remote site, the design provides a consistent set of features and functionality for network access, which helps improve user satisfaction and productivity while also reducing operational expense. The ESP Campus delivers a consistent and scalable methodology of building the network, improving overall usable network bandwidth and resilience, while making it easier to deploy, maintain, and troubleshoot.

What’s New in This Version

The following changes were made since Aruba last published this guide:

- This is a new guide.

www.arubanetworks.com

3333 Scott Blvd. Santa Clara, CA 95054
1.844.472.2782 | T: 1.408.227.4500 | FAX: 1.408.227.4550

ESP-CPDS-21A-1 06/21