Aruba Mobility Controllers and Deployment Models Validated Reference Design

Version 5.0

ARUBA networks

Solution Guide

Contents

Chapter 1

Introduction

Aruba Reference Architectures

The Aruba Validated Reference Design (VRD) series is a collection of technology deployment guides that include descriptions of Aruba technology, recommendations for product selections, network design decisions, configuration procedures, and best practices for deployment. Together these guides comprise a reference model for common customer deployment scenarios. Each Aruba VRD network design has been constructed in a lab environment and thoroughly tested by Aruba engineers. Our customers use these proven designs to rapidly deploy Aruba solutions in production with the assurance that they will perform and scale as expected.

The VRD in this guide discusses deployment models and mobility controller selection. This guide describes:

- Operating modes for the mobility controller
- Licensing
- Forwarding modes
- Logical and physical deployment
- Redundancy
- How to select the appropriate mobility controller based on scalability requirements

This guide will help you understand the capabilities and options you have when deploying an Aruba Mobility Controller. Other guides in the series will build specific deployments using the information in this guide.

Reference Documents

This VRD is the first in a series of new guides in a revised format, and it is intended to provide the most current information about the mobility controller technology and deployment models introduced in ArubaOS 5.0. Other guides in the series will cover 802.11n and AP deployments, user roles and firewall policies, outdoor deployments, and configurations for specific deployments such as campus or remote networks. This guide is a foundation-level guide, and therefore it will not cover the configuration of the Aruba system. Instead, it is intended to provide the baseline knowledge that a wireless engineer must use to deploy an architecture that is based on the mobility controller.

For specific deployment configuration details, see the 3.X series of VRDs on the Aruba website at http://www.arubanetworks.com/vrd. The existing VRDs will be updated to follow this new format.

The complete suite of Aruba technical documentation is available for download from the Aruba support site. These documents present complete, detailed feature and functionality explanations beyond the scope of the VRD series. The Aruba support site is located at: https://support.arubanetworks.com/.

For more training on Aruba products or to learn about Aruba certifications, visit our training and certification page on our public website. This page contains links to class descriptions, calendars, and test descriptions: http://www.arubanetworks.com/training.php

Aruba hosts a user forum site and user meetings called Airheads. The forum contains discussions of deployments, products, and troubleshooting tips. Airheads is an invaluable resource that allows network administrators to interact with each other and Aruba experts. Announcements for Airheads meetings are also available on the site: http://airheads.arubanetworks.com/

The VRD series assumes a working knowledge of Wi-Fi®, and more specifically dependent AP-based architectures. For more information about wireless technology fundamentals, visit the Certified Wireless Network Professional (CWNP) site at http://www.cwnp.com

Chapter 2

Summary of Recommendations

This VRD explains the recommendations that Aruba provides for deploying a scalable, mobility controller-based wireless LAN (WLAN) system. The details of the system occasionally require lengthy explanations and diagrams to sufficiently convey the details of the solution. This section provides an overview of the Aruba recommendations for mobility controller deployments, which are explained in detail in the following sections.

Campus Deployments

Campus-based deployments are networks that require more than a single controller to cover a contiguous space. Examples of campus-based deployments are corporate campuses, large hospitals, and higher-education campuses. In these deployments, the WLAN is often the primary access method for the network, and it is typically used by multiple classes of users and devices.

Figure 1 *Typical campus deployment with redundancy*

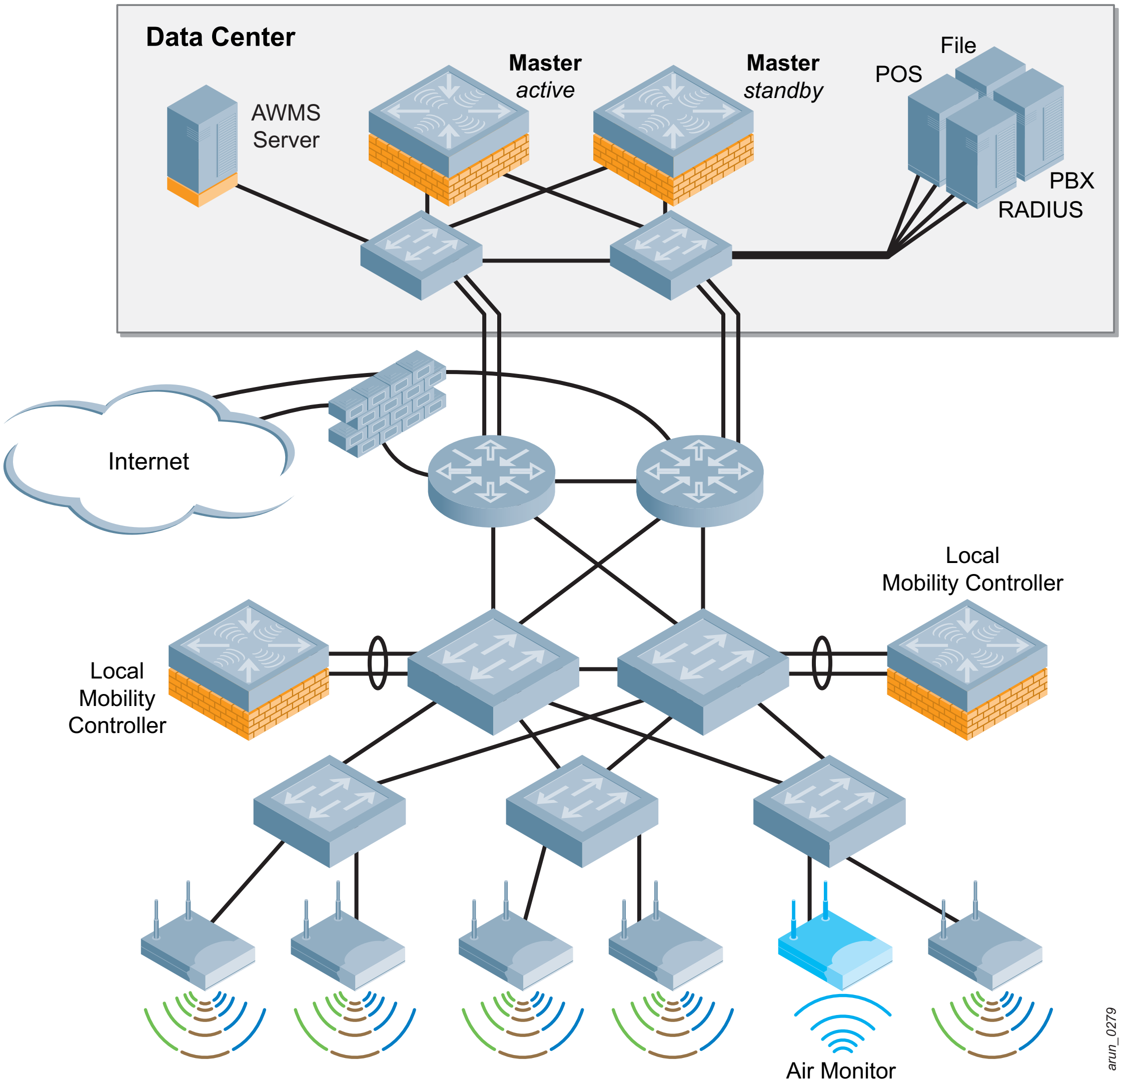

The following table summarizes the recommendations for campus-based deployments.

Table 1 *Campus Deployment Recommendations*

Feature or Function	Recommendation
Master/local or all masters	Use a master/local configuration.
Physical location of the master mobility controller	Put the master mobility controller in the network data center.
Master mobility controller redundancy	Use master redundancy with a standby master.
Physical location of the local mobility controller	Put the local mobility controller in the distribution layer or the data center, depending on the traffic patterns in the organization and the location of the Internet drop.
Local mobility controller redundancy	Use active-active with two mobility controllers and two instances of the Virtual Router Redundancy Protocol (VRRP). Each mobility controller should be the primary on one of the VRRP instances, and back-up for the other. Half of the APs should terminate on each of the two VRRP instances.
Local mobility controller load	AP and user load to 40% capacity, which results in 80% load in a failure condition.
Control plane security (CPsec)	Off by default, but when network administrators plan the network, ensure that future CPsec operation is possible.
AirWave Wireless Management Suite (AWMS)	Deploy AWMS for management, configuration, reporting, and troubleshooting. Deploy AWMS in the network data center.
Licenses	● AP capacity equal to the number of campus access points (CAPs) and air monitors (AMs) supported during a failover scenario. ● Policy Enforcement Firewall – Next Generation (PEF-NG) for user roles, firewall policy equal to AP capacity supported during a failover scenario. ● Wireless Intrusion Prevention (WIP) for advanced wireless intrusion protection and rogue detection, equal to AP capacity supported during a failover scenario. ● xSec for Federal and other Government deployments.

Remote Deployments

For deployments that cover remote access, two solutions exist. The remote access point (RAP) provides secure, clientless access to the small branch, home office, and fixed telecommuter. These deployments are typically characterized by the need for multiple network components, such as Voice over IP (VoIP) phones, wireless printers, and local disk storage. For the highly mobile user, the Aruba Virtual Internet Access (VIA) IPsec client provides seamless connectivity without the need for local infrastructure. The VIA client works over Wi-Fi, Ethernet, and cellular connections.

Figure 2 *Remote deployment with RAPs and VIA clients*

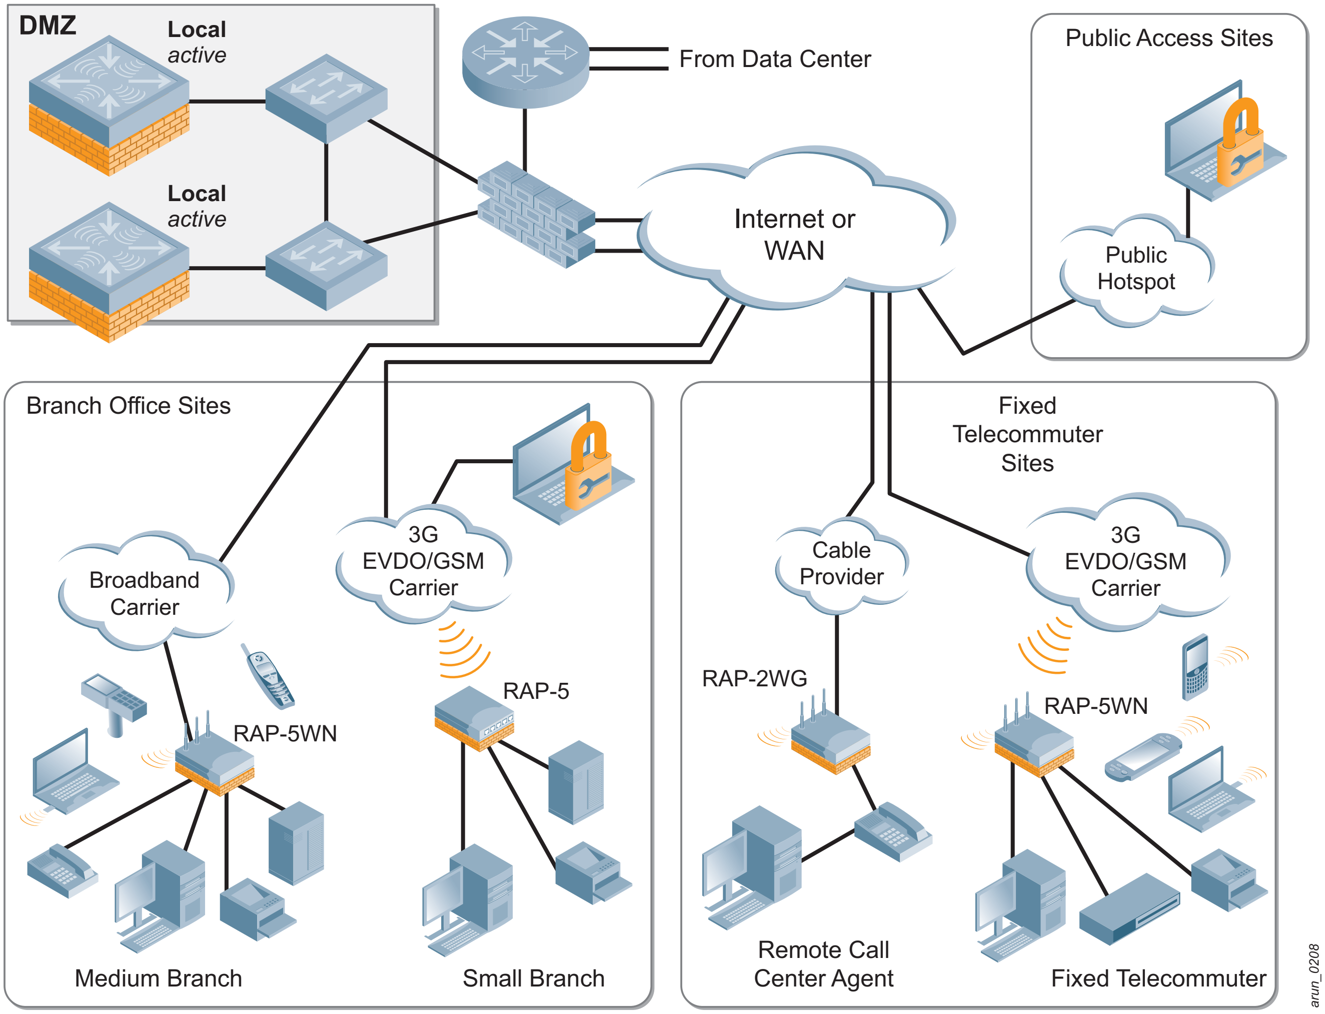

The following table summarizes the recommendations for RAP and VIA-based deployments.

Table 2 *Remote Deployment Recommendations*

Feature or Function	Recommendation
Master/local or all masters	Use a master/local configuration.
Physical location of the master mobility controller	**Network data center:** Put the master mobility controller in the network data center if the Aruba deployment is pervasive. **DMZ:** Put the master mobility controller in the DMZ if Aruba is used only for remote networking.
Master mobility controller redundancy	Use master redundancy with a standby master.
Physical location of the local mobility controller	Put the local mobility controller in the network DMZ.
Local mobility controller redundancy	Use active-active, with two mobility controllers and two instances of the VRRP. Each mobility controller should be the primary on one of the VRRP instances, and back-up for the other. Half of the APs should terminate on each of the two VRRP instances.
Local mobility controller load	AP and user load to 40% capacity, which results in 80% load in a failure condition.
VIA agents	Use the VIA agent on laptops that need to communicate over untrusted connections where a RAP deployment is not feasible.
AirWave Wireless Management Suite	Deploy AWMS for management, configuration, reporting, and troubleshooting. Deploy AWMS in the network data center.
Licenses	● AP capacity equal to the number of RAPs supported during a failover scenario. ● PEF-NG for user roles, firewall policy, equal to the number of RAPs supported during a failover scenario. ● PEF-V for VIA client firewall policy. ● WIP for advanced wireless intrusion protection and rogue detection, equal to the number of RAPs supported during a failover scenario.

Branch Offices

The branch office is an extension of a larger organization. Typically, the branch office has a data center located at a remote site where additional services are provided. The branch office is characterized by being large enough to require multiple APs to service local clients. Often the branch office has a requirement for survivability in the event of a WAN outage.

Figure 3 *Branch office controller deployment with guest access*

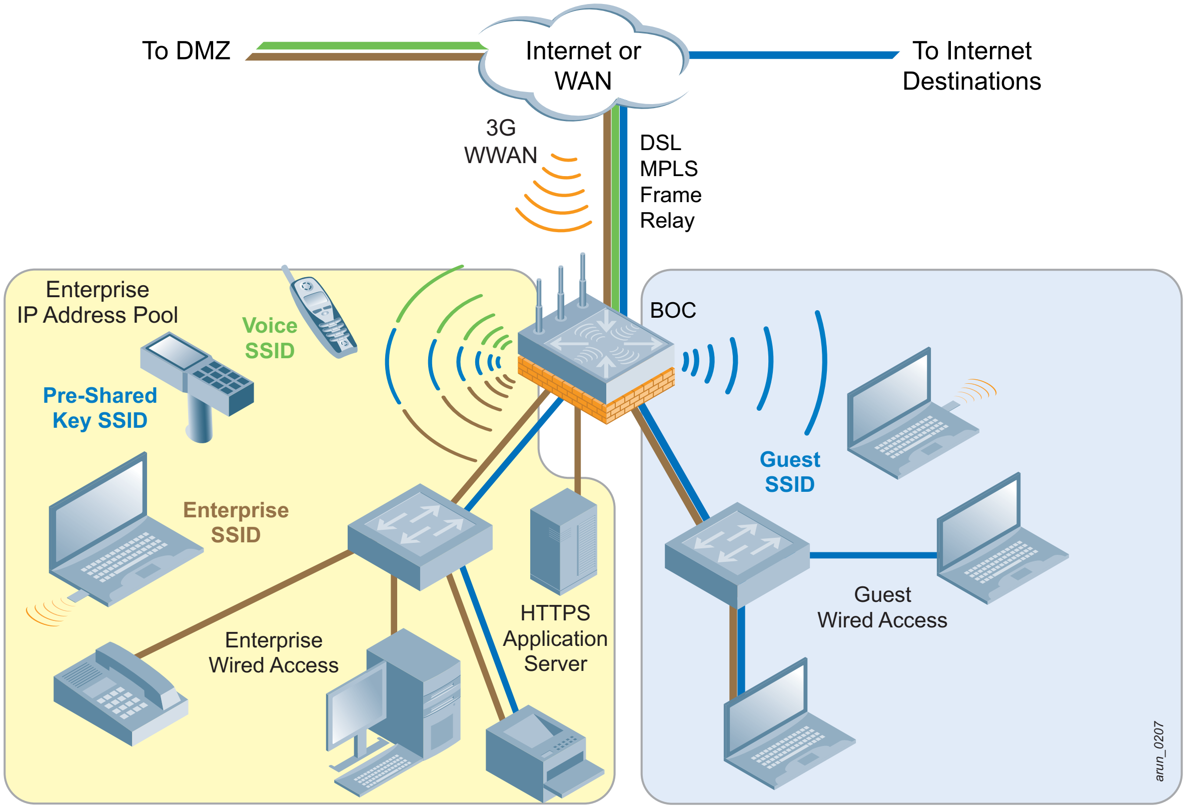

The following table summarizes the recommendations for branch office deployments, with the assumption that a central corporate site exists elsewhere.

Table 3 *Branch Office Recommendations*

Feature or Function	Recommendation
Master/local or all masters	Use an all masters configuration.
Location of the master mobility controller	Put the master mobility controller in the networking closet.
Master mobility controller redundancy	Use master redundancy using a standby-master, or N+1 for larger branches with private WAN connections.
Location of the local mobility controller	N/A
Local mobility controller redundancy	N/A
Local mobility controller load	N/A
VIA agents	Use the VIA agent on laptops that need to communicate over untrusted connections where a RAP deployment is not feasible.
AirWave Wireless Management Suite	Deploy AWMS for management, configuration, reporting, and troubleshooting. Deploy AWMS in the network data center.
Licenses	● AP capacity equal to the number of RAPs supported during a failover scenario. ● PEF-NG for user roles, firewall policy, equal to the number of RAPs supported during a failover scenario. ● PEF-V for VIA client firewall policy. ● WIP for advanced wireless intrusion protection and rogue detection, equal to the number of RAPs supported during a failover scenario.

Small Offices with Single Sites

In some cases, an organization is small enough that a single mobility controller is sufficient to handle the networking needs of the organization. Typical cases include doctors offices, law firms, design firms, and architects offices. Typically the organization either has no other site, or they are few in number and can be handled by RAPs and VIA clients that terminate on the same mobility controller.

Figure 4 *Small branch deployment with a mobility controller onsite*

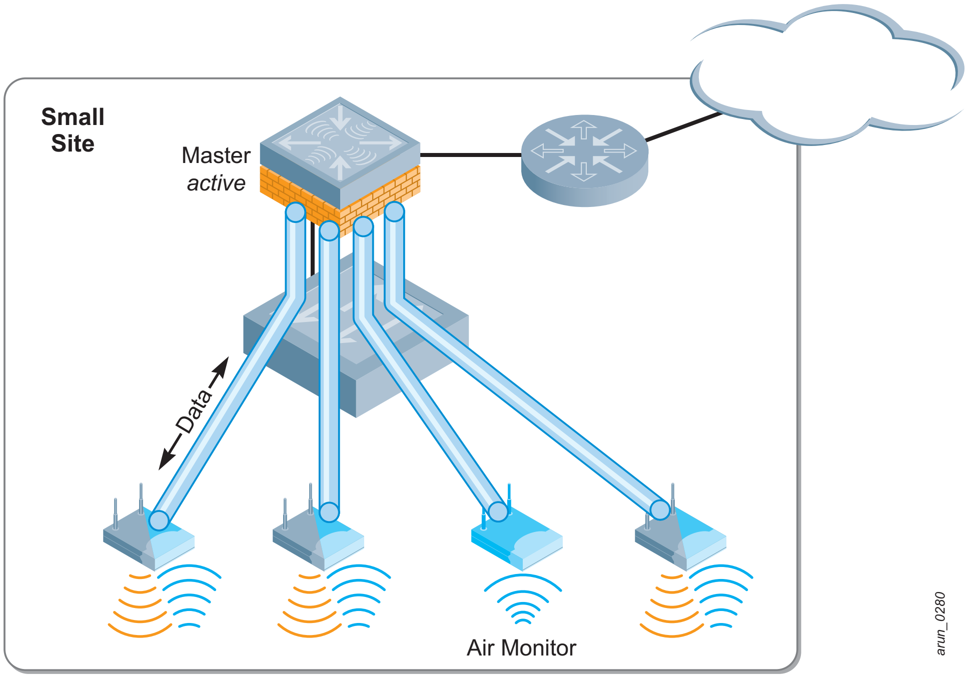

The following table summarizes the recommendations for small organizations.

Table 4 *Small Office Deployment Recommendations*

Feature or Function	Recommendation
Master/local or all masters	Use an all masters configuration.
Location of the master mobility controller	Put the master mobility controller in the networking closet.
Master mobility controller redundancy	Use master redundancy, using a standby-master if possible.
Location of the local mobility controller	N/A
Local mobility controller redundancy	N/A
Local mobility controller load	N/A
VIA agents	Use VIA agents on laptops that need to communicate over untrusted connections where a RAP deployment is not feasible.
AirWave Wireless Management Suite (AWMS)	If possible, deploy AWMS for management, configuration, reporting, and troubleshooting. Deploy AWMS in the network data center.

Table 4 *Small Office Deployment Recommendations (Continued)*

Feature or Function	Recommendation
Licenses	● AP capacity equal to the number of RAPs supported during a failover scenario. ● PEF-NG for user roles, firewall policy, equal to the number of RAPs supported during a failover scenario. ● PEF-V for VIA client firewall policy. ● WIP for advanced wireless intrusion protection and rogue detection, equal to the number of RAPs supported during a failover scenario.

Chapter 3

Understanding the Aruba Mobility Controller and AWMS

The Aruba Mobility Controller is the heart of the Aruba dependent access point (AP) WLAN architecture. The mobility controller is responsible for many of the operations that traditionally would be handled by the AP in an autonomous AP deployment. The mobility controller acts as a command-and-control point for the network as a whole. The mobility controller operates as an access network with APs and as a wireless intrusion prevention (WIP) system with dedicated AMs that perform scanning and containment operations on the WLAN.

The Aruba Mobility Controller goes far beyond managing dependent APs. The mobility controller is capable of fulfilling many of the roles that traditionally were handled by dedicated appliances. The functionality that the mobility controller provides includes:

- Acting as a user-based stateful firewall
- Terminating user-encrypted sessions from wireless devices
- Performing Layer 2 switching and Layer 3 routing
- Providing clientless Layer 3 mobility
- Acting as an IPsec virtual private network (VPN) concentrator for site-to-site and client-based VPNs
- Providing certificate-based IPsec security to protect control channel information
- Terminating Internet-based RAPs
- Providing wired firewall services
- Performing user authentication with 802.1X and captive portal authentication, among others
- Providing guest access and provisioning services
- Providing advanced RF services with Adaptive Radio Management (ARM) and spectrum analysis
- Providing location services and RF coverage "heat maps" of the deployment
- Performing rogue detection and containment
- Providing self-contained management by way of a master/local hierarchy with one controller pushing configuration to other mobility controllers to reduce administrative overhead
- Delivering AP software updates automatically when the mobility controller is upgraded

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

For larger or more distributed organizations, the Aruba AirWave Wireless Management Suite (AWMS) provides overall configuration and reporting capabilities to complement the management tools that are built into the controller. Features include configuration across multiple master/local clusters, long-term reporting, location services, and advanced rogue AP detection tools in a highly scalable solution.

The Aruba system has a logical four-tier operating model that consists of four layers: management, network services, aggregation, and network access. Controllers operate at the network services and aggregation layers, APs operate at the network access layer, and AWMS operates at the management level. This model is used throughout the VRD program to describe the logical deployment of the Aruba solution.

Figure 5 *Logical four-tier operating model*

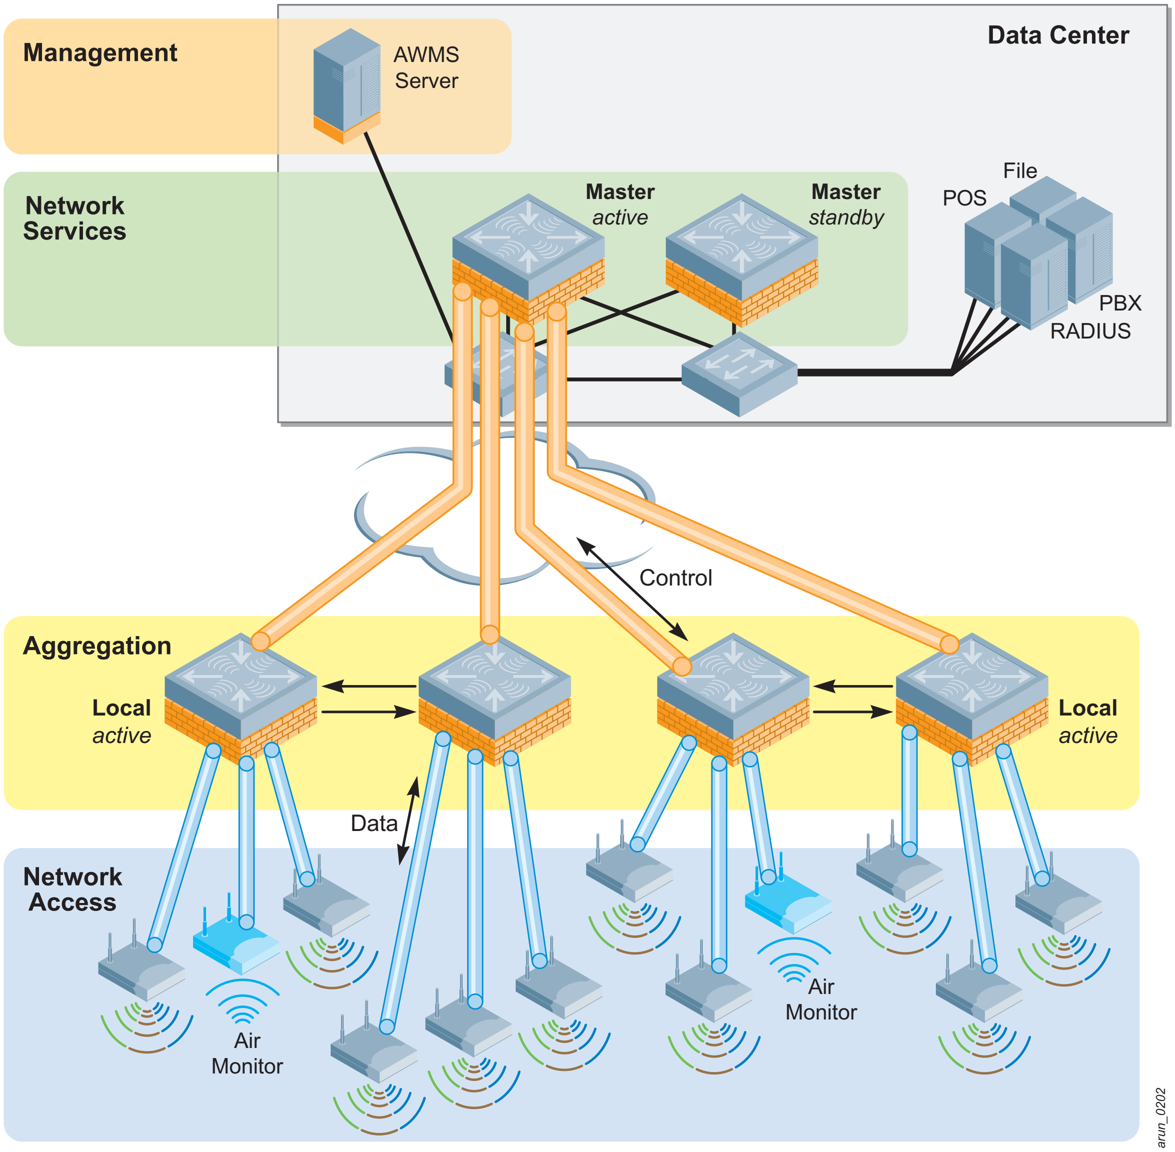

Controller Overview

The mobility controllers are available in modular chassis models and as network appliances that scale to meet the needs of the largest organizations. This section introduces the current generation of mobility controllers.

Aruba 6000 Chassis and M3 Blade

Figure 6 *Aruba 6000 Chassis with four M3 Mobility Controller Blades*

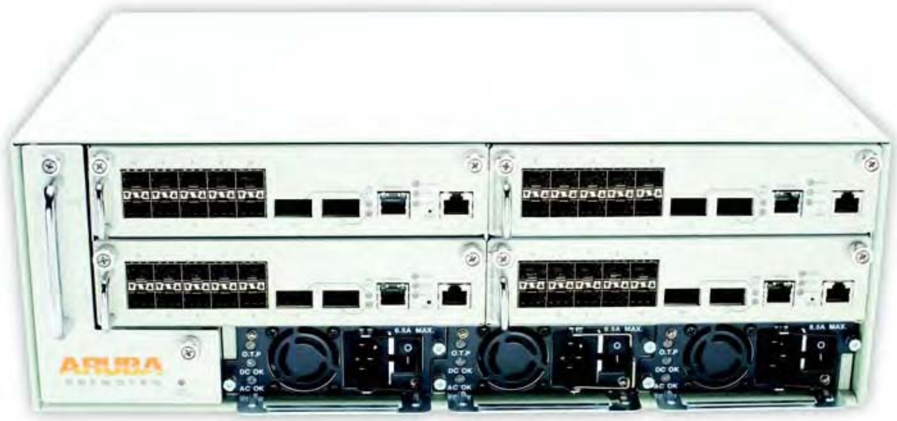

The Aruba 6000 Chassis is designed to address a wide range of wireless and wired network mobility, security, and remote networking requirements for corporate headquarters and large campus deployments. The Aruba 6000 is easy to install without disrupting the existing wired network. Edge services are virtualized and implemented in cost-effective APs at the network edge. The APs move user traffic to data center controllers through secure IP tunnels over a public or private transport network. The Aruba 6000 runs the ArubaOS operating system and comes standard with advanced authentication, encryption, wireless radio management, secure enterprise wireless mesh, and Layer 2 and Layer 3 networking features. Optional ArubaOS software modules deliver additional functionality, including policy enforcement firewall, VPN server, remote access gateway, and WIP.

Aruba 3000 Series

Figure 7 *Aruba 3000 Series Mobility Controllers*

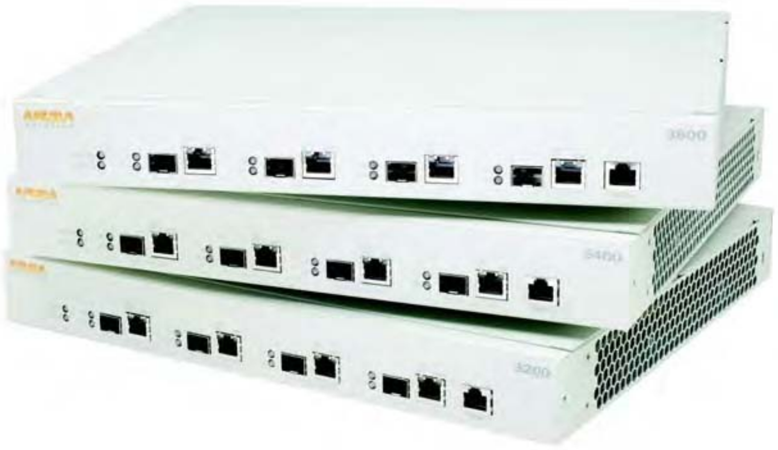

The Aruba 3200 Mobility Controller is designed for the small and branch offices, and the Aruba 3400 and Aruba 3600 Mobility Controllers are designed for medium and large enterprise or dense office deployments. Edge services are virtualized and implemented in cost-effective APs at the network edge. The APs move user traffic to data center controllers through secure IP tunnels over a public or private transport network. The Aruba 3000 Series runs the ArubaOS operating system and comes standard with advanced authentication, encryption, wireless radio management, secure enterprise wireless mesh, and Layer 2 and Layer 3 networking features. Optional ArubaOS software modules deliver additional functionality, including policy enforcement firewall, VPN server, remote access gateway, and WIP.

600 Series

Figure 8 *651 Controller Branch Office Controller with Integrated AP*

The Aruba 600 Series of controllers are integral components of the Aruba distributed enterprise network, which uses central controllers in the data center to manage complex and processing-intensive management and security functions. Edge services are virtualized and implemented in cost-effective edge devices, which move user traffic to data center controllers through secure IP tunnels over a public or private transport network. The Aruba 600 Series provides local wireless services when used in conjunction with any Aruba AP, or they can operate as wired-only devices. The controllers also include print server and network-attached storage capabilities to enable local network printing and mass storage.

Next Steps

The next three sections discuss the deployment operating modes and deployment options for the system components. The first section describes the different operating modes of the Aruba Mobility Controllers and explains the functions of the master and local controller. The second section explains the advantages of adding AWMS to the WLAN to increase management capabilities. Finally, the third section describes the various deployment models that allow Aruba WLANs to scale from the largest campus to the smallest home office.

Understanding Mobility Controller Master/Local Model

All Aruba Mobility Controllers are capable of assuming two operating roles in the system: master or local mobility controller. This hierarchy allows organizations to build scalable WLAN networks without requiring additional management platforms when the network is contained to a single master/local cluster.

The master mobility controller is the central point of coordination and configuration of the network. The master mobility controller processes all wireless security events and sends policy-based configuration to the local mobility controllers. The local mobility controllers manage the CAPs, air monitors (AMs), RAPs, VPN clients, and devices attached to the WLAN. APs connect directly to the local mobility controller over an IP-based network, and in most deployments, all traffic from devices is sent to the local mobility controller for processing.

Understanding the Master Mobility Controller

The role of the master mobility controller in an Aruba WLAN is to provide a single point of policy configuration and coordination for the WLAN in smaller deployments. The master mobility controller can receive configuration and coordination information from the AWMS for larger or more distributed deployments. A typical master/local cluster consists of one master mobility controller and one or more local mobility controllers. However, in smaller deployments the master also can handle all functions of the local. The communication channel between the master and locals is protected using IPsec. The master mobility controller typically does not terminate any clients or APs directly.

Figure 9 *Network services layer*

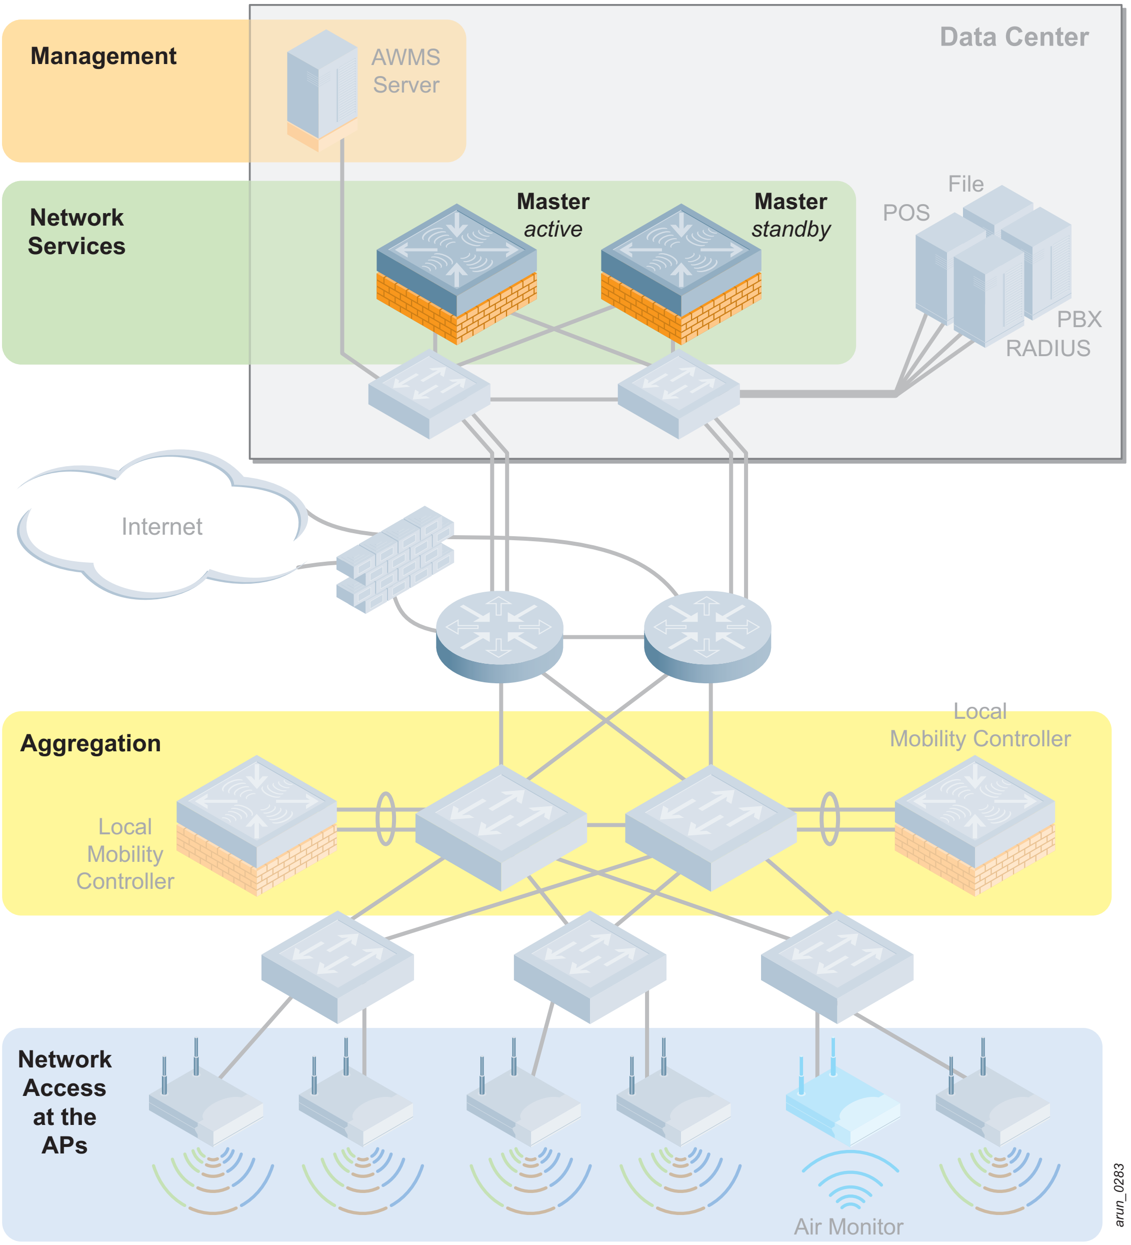

Master mobility controllers are responsible for the following functions in the WLAN:

- **Policy configuration:** Configuration in the Aruba solution is split between policy and local configurations. Local configuration deals with things such as physical interfaces, IP networking, and VLANs. Policy configuration is centered on the operation of APs and users, including AP settings such as the SSID name, encryption, regulatory domain, channel, power, and ARM settings. Policy configuration extends beyond APs and also covers user authentication, firewall policy, mobility domains (IP mobility), IPsec settings, and system management. The policy is pushed to all local mobility controllers in the form of profiles, and profiles combine to create the configuration for the dependant APs.
- **AP white lists:** Two types of white lists exist in the system, one for RAPs and one for CAPs that use CPsec. These lists determine which APs can connect to the mobility controllers and in which method of operation. These lists give organizations additional levels of control. Unauthorized devices are prevented from connecting to the network, and provisioning of the system is simplified.
- **WIP coordination:** Wireless intrusion prevention (WIP) activities involve looking for rogue (unauthorized) APs and monitoring for attacks on the WLAN infrastructure or clients. The master controller processes all data collected by Aruba Wi-Fi sensors known as AMs. Instructions to disable a rogue AP or blacklist a client from the network are issued through the master mobility controller.
- **Valid AP list:** Separate from the concept of the white list, all of the legitimate APs operating on the WLAN must also be known to all mobility controllers in the network, and they must be marked as being valid APs. This list prevents valid APs from being falsely flagged as rogue APs when APs that are attached to two different local mobility controllers are close enough to hear each other's transmissions on a single site. This list also helps ARM to differentiate between APs that generate noise (third-party or not under control of the mobility controller) and those that are part of the valid coverage area, but are generating co-channel interference. Unlike traditional wireless intrusion detection system (WIDS) solutions, the master controller automatically generates the valid AP list without network administrator intervention. All Aruba APs are automatically learned and added to the list, but valid third-party APs must be added manually. If more than one master/local cluster exists, AirWave should be deployed to coordinate APs between clusters.
- **RF visualization:** Unlike traditional wired networks, RF requires additional tools to visualize network coverage. The Aruba RF visualization tools provide a real-time view of the network coverage. The entire network or coverage in a particular band, channel, or speed can be selected. This information is based on the AP channel and power settings and the data collected from AMs and APs listening to transmissions during their scanning periods. This information provides a realistic picture of the actual RF coverage as opposed to predictive planning tools. Predictive tools cannot provide anything more than a model of what should occur in the network based on estimates using AP placement and wall materials.
- **Location:** Much like visualization, locating users in the WLAN is more difficult with mobile clients and IP mobility. The WLAN is not only serving clients but is also scanning off channel, so it is possible to triangulate users and rogue devices to within a small area. This information is displayed on the master and allows for devices to be located quickly, which is critically important for physical security and advanced services such as E911 calling.
- **Initial AP configuration:** When an AP first boots up, it contacts its master mobility controller to receive the configuration generated by the master. The AP can use DNS lookup, DHCP option 43, broadcast, multicast, or static configurations to locate the master, which creates a "plug-and-play" experience when APs are added to the network. The master compares the AP name and determines its group assignment, and then redirects that AP to the proper local mobility controller. The most common methods are DNS and DHCP options, which allow changes to propagate quickly through the system and eliminate the need for Layer 2 connectivity to the master.
- **Control plane security:** When CPsec is enabled, the master mobility controller generates the self-signed certificate and acts as the certificate authority (CA) for the network. The master mobility controller issues certificates to all local mobility controllers in the network as well. If more than one

master exists in the network, the network administrator assigns a single master as the trust anchor for that network. The trust anchor issues certificates to the other master controllers in the network.

- **Authentication and roles:** User authentication methods and role assignments are created on the master mobility controller and then propagated to local mobility controllers throughout the network. A database exists to authenticate users in small deployments or for guest access credentials that can be leveraged by all the mobility controllers in the network. Additionally, the master can also proxy requests for the network to a RADIUS or LDAP server.

Understanding the Local Mobility Controller

The local mobility controller controls its logically attached APs and handles user sessions on the network. The locals process the majority of the traffic on the network. When the locals terminate CAPs, the locals are typically deployed either in the distribution layer or network data center, depending on the distribution of traffic in the enterprise. In the case of RAPs, Branch Office Controllers (BOCs), and VIA agents, the locals are typically located in the network DMZ.

Figure 10 *Aggregation layer*

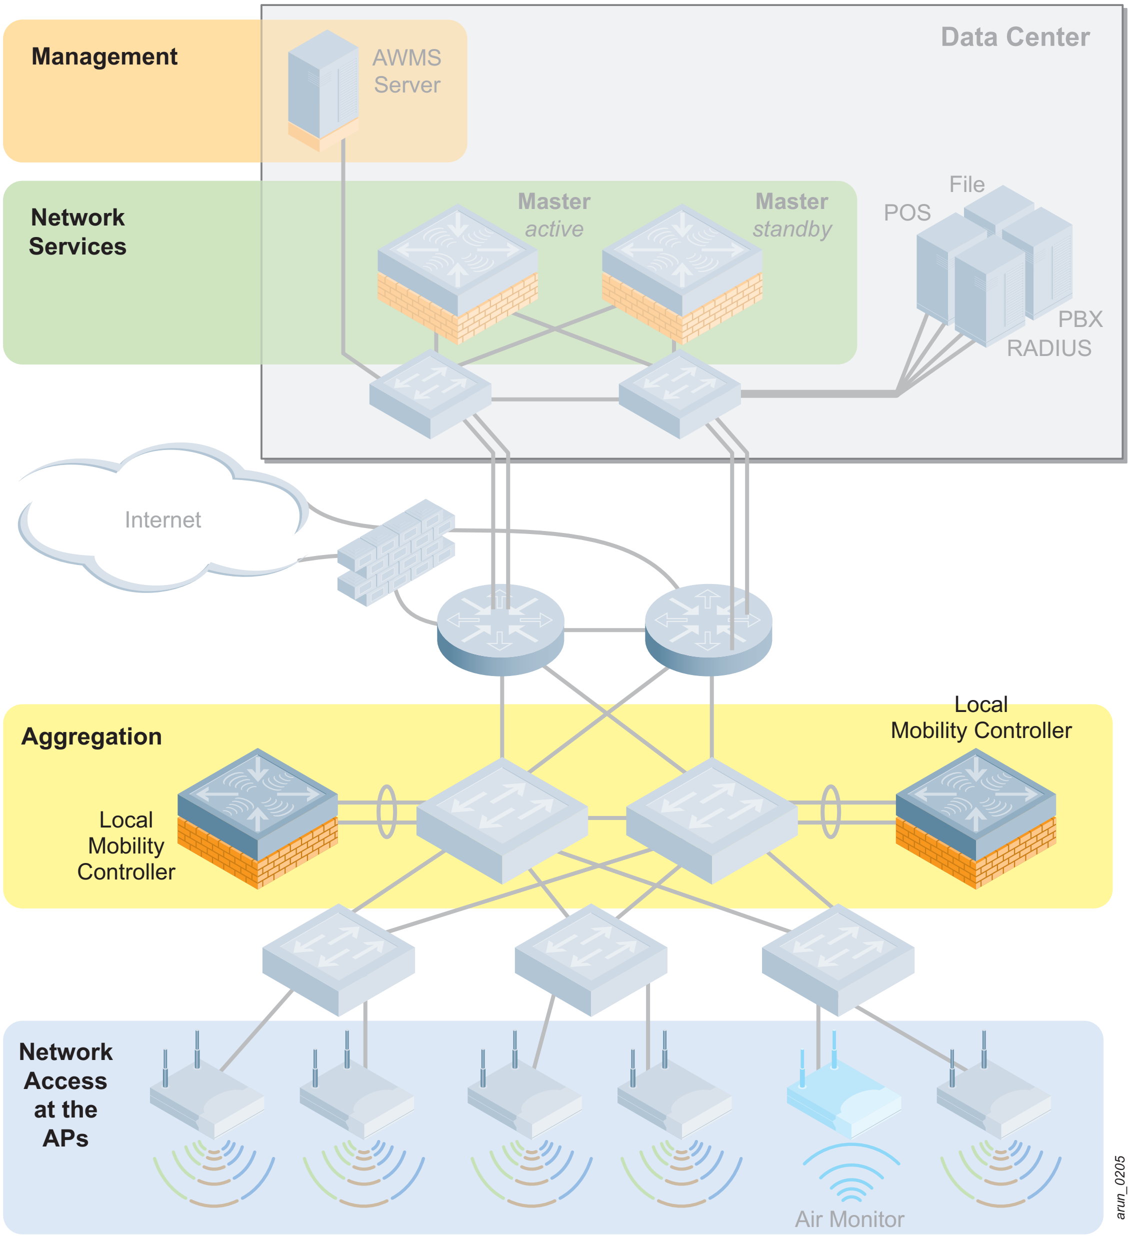

Local mobility controllers are responsible for the following functions in the WLAN:

- **AP and AM configuration and software updates:** All Aruba APs are dependent APs, which means they do not in most instances store configuration settings in the way that a traditional autonomous AP would. Instead, at boot time each AP downloads its current configuration from the local mobility controller. When changes are made in the system configuration, they are automatically pushed to all APs. Whenever an AP boots, it will always have the current configuration, and changes are reflected immediately throughout the network. When the software on the mobility controller is updated, the APs automatically download a new image from the mobility controller and upgrade themselves. This software check, like the configuration download, is part of the AP boot process, and it insures that each AP has the current operating image and configuration without user intervention.
- **AP and AM termination:** The local mobility controller is the location where the AP terminates after it contacts the master mobility controller. The local upgrades the APs operating system as necessary, and it pushes the APs configuration to it. In the majority of campus deployments, all user traffic is sent from the AP to the local mobility controller and the local handles the decryption of the wireless frame. AMs also terminate on the local, but their traffic is sent directly to the master mobility controller to handle WIP coordination between AMs and APs. In an Aruba system, user traffic is sent over GRE tunnels to the local mobility controller. Control traffic also uses a protocol called PAPI, but if CPsec is enabled, control traffic uses an IPsec tunnel.
- **User session termination:** An Aruba network is focused on the client devices, which are commonly referred to as users. User sessions are any information transmitted from a client device across the WLAN, from actual human users on a wireless device to wireless IP cameras to medical equipment and scanner guns. Every user in an Aruba system is identified when they authenticate to the system (by WLAN, IPsec, or wired with captive portal), and their login information is used to place the user in the appropriate role based on that login. The user's role defines what that user is allowed to do on the network. This definition is enforced by an ICSA [1]certified stateful firewall, with a role-based policy applied to every user. By applying the policy to the individual users, Aruba eliminates the need to focus on ports as a security mechanism. Instead security moves to the user level, which allows multiple users to share a single AP or wired port and receive the appropriate level of security based on whom they are, not where they connect.
- **ARM assignments and load balancing:** Aruba ARM controls aspects of AP and client performance. All WLANs operate in unlicensed space, so the chance that something will interfere is very high. Aruba has developed a system to automatically work around interference and help clients have a better operating experience. These features include automatically tuning the WLAN by automatically configuring AP power and channel settings, as well as scanning for better settings and avoiding interference. ARM also handles AP load balancing and co-channel interference from other APs and clients, as well as airtime fairness to ensure that slower speed clients do not bring down the throughput of higher-speed clients. Using band steering, when the system detects a client is capable of operating on the 5 GHz band (the majority of modern clients), the system automatically attempts to steer that client to the cleaner band.
- **WIP enforcement and blacklisting:** While the master handles the processing of WIP information, the local directs the actions of the AMs for enforcement of WIP policy. Enforcement can take different shapes, including containing rogue APs by performing denial-of-service (DoS) attacks wirelessly, ARP cache poisoning on the wire, shielding valid clients from connecting to rogue APs, and blacklisting clients so that they are unable to attach to the WLAN.

[1] ICSA labs provides vendor neutral testing of products and certifies them in compliance with a set of common tests and criteria. ICSA is on the web at http://www.icsalabs.com/

- **CPsec AP certification:** When CPsec is enabled in the WLAN, the AP and local mobility controller establish an IPsec tunnel between the two devices using certificates. The local mobility controller is responsible for issuing these certificates, and in some instances adding APs to the white list. When the AP boots up and tries to contact the local mobility controller, the certificates are used to build an IPsec tunnel between the devices.
- **Mobility:** Supports both Layer 2 (VLAN) mobility as well as Layer 3 (IP) mobility, allowing users to seamlessly roam between APs on different mobility controllers without session interruption. This is a key component to support VoIP sessions, where sessions must be preserved.
- **Quality of service (QoS):** The mobility controllers support QoS on the wired and wireless side. This support includes translating DiffServ and ToS bits set on packets into Wi-Fi Multimedia™ (WMM®) markings and back. The Aruba policy enforcement firewall (PEF) also allows the administrator to mark packets with the appropriate level of QoS, and to change markings on packets entering the system.

Understanding the Role of the AirWave Wireless Management Suite

The management framework built in to the master mobility controller was developed to allow small-scale networks to operate without the need of a dedicated management platform. When the system was first developed, Wi-Fi networks were convenience networks that were used only in meeting rooms and usually by guests. As Wi-Fi moved from providing only guest access to being the primary connection medium, the mission-critical nature demanded new tools to manage the network.

Figure 11 *Management layer*

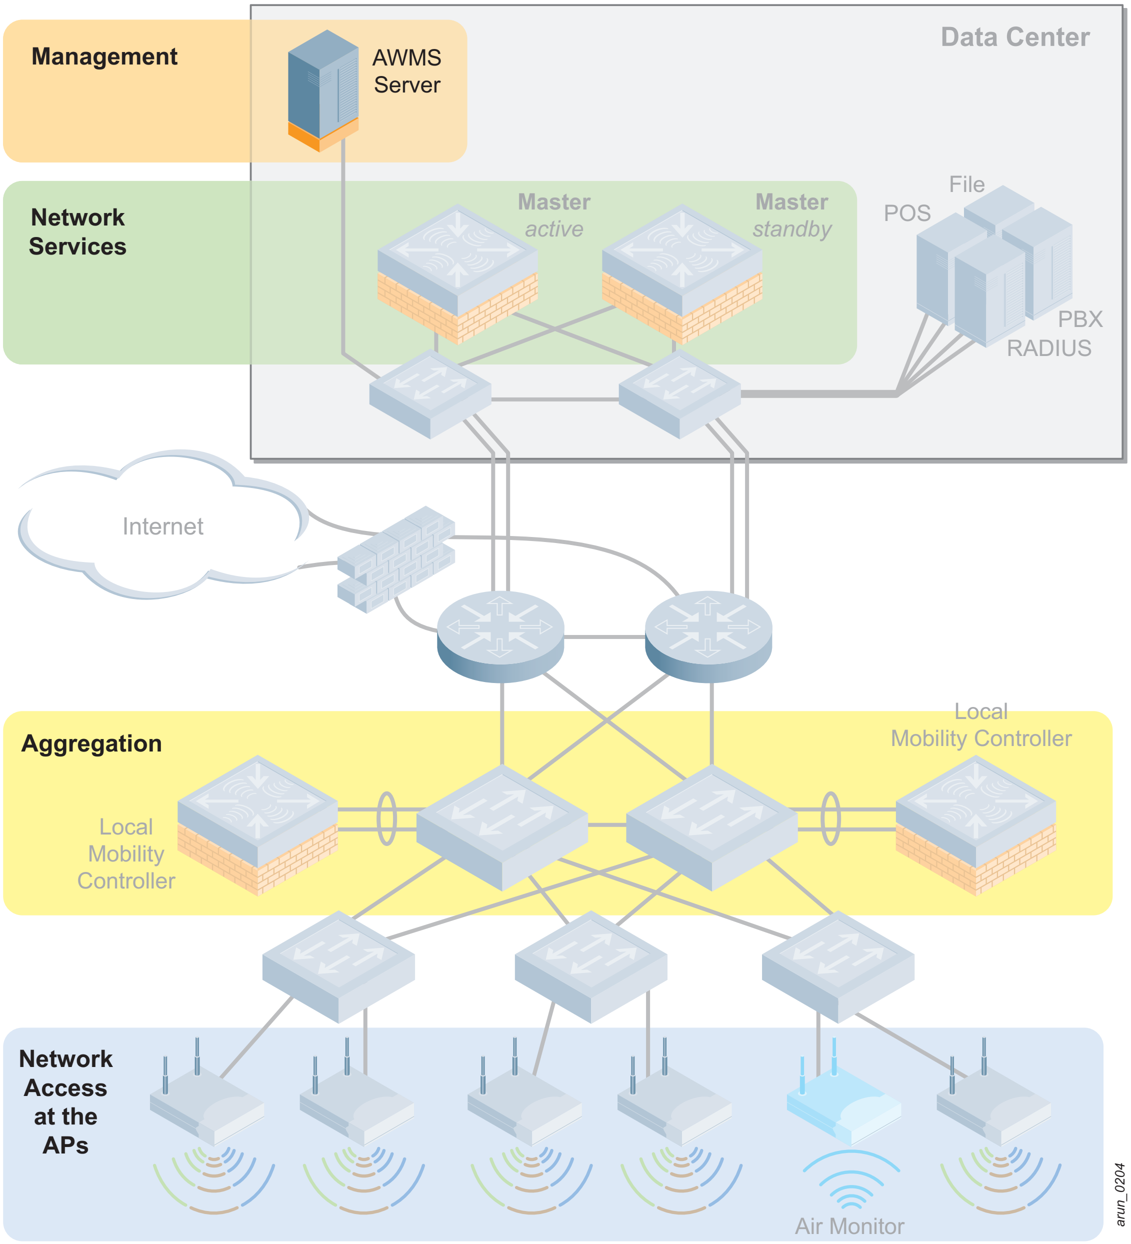

The AirWave Wireless Management Suite (AWMS) fills this role by providing best-in-breed network management. AWMS enhances the features that are already delivered by the master controller and allows the WLAN to scale to much larger deployments. AirWave increases the functionality and usability of the Aruba Mobility Controller control plane in the following ways:

- **Historical data:** The largest Aruba Mobility Controller contains 512 MB of flash memory, and an AirWave appliance contains multiple hard drives with 100+ GB of capacity. The Aruba Mobility Controller can typically store 1-4 days of data, and an AirWave appliance can store 1+ year.
- **Scalability:** A master Aruba Mobility Controller can effectively manage up to 50,000 objects including APs, clients, and local mobility controllers. To avoid overloading a master mobility controller, real-time user information is also limited in larger networks. The controller polls for real-time data on-demand rather than continuously collecting it. The AirWave Master Console™ can be used to manage over 75,000 APs, hundreds of controllers, and hundreds of thousands of users. Data is collected continuously and in real time so that problems can be diagnosed after the fact by using historical data.
- **Reporting:** To assess performance and network capacity, it is not enough to look at real-time network data in isolation. Trend data from months or years can help to determine how current conditions compare to previous levels and how changing usage patterns may impact network performance. AWMS provides up to 2 years of actionable information, including network performance data and user roaming patterns. AWMS also provides the detailed capacity reports that are needed to plan for adequate coverage across the current network, network expansion, or 802.11n upgrades. Reports can be scheduled to run automatically and are emailed out, which eliminates the need to manually generate reports or log into the server to check status.
- **Alerting:** Automated alerts can be sent using SNMP, syslog, or email, and the alerts can be configured to cover the entire system or only portions of the network. Alerting is available for a range of issues including radio state, user bandwidth, and rogue AP detection.
- **Advanced RF tools:** AWMS can provide real-time views of the RF coverage in an organization and provide “heat maps” of that coverage on floor plans where APs and AMs are deployed. Additionally, these same maps can be used to locate client stations, rogue APs, and 802.11 RFID tags. Additional tools make it possible to plan network deployments using VisualRF Plan to do AP deployment planning software, without the need for more expensive manual surveys in many cases. This planning is backed up by VisualRF, which delivers a real-time view of the network as it operates.
- **Managing multiple master/local clusters:** An AWMS can monitor and control multiple master Aruba Mobility Controllers and their associated local mobility controllers. Configuration auditing and reporting is also available to ensure that all devices are configured as expected, and that changes to the system are recorded. In addition, software can be pushed to the mobility controllers along with scheduled reboots to upgrade the system.
- **Executive and help desk monitoring views:** The ArubaOS WebUI is primarily intended for network engineers to conduct configuration and monitoring. An AirWave Management Suite provides multiple customized interfaces appropriate to different roles in the organization. AWMS also allows for multiple user roles with different permissions on the system, as well as the ability to group devices to provide permissions to a subset of the network. Open APIs allow for integration into existing network management and trouble ticket systems.
- **Management of third-party devices** – The AWMS is the premiere multivendor wireless management suite on the market. Many large organizations have a mixture of wireless equipment, new and old, and are slowly migrating from one system to the other. The AWMS allows network managers to transition from legacy platforms to the Aruba solution using a single interface for management.
- **Wired network rogue AP scans:** The AirWave RAPIDS module can scan networks for rogue APs even when no wireless sensors are present. When the RAPIDS module is combined with the ArubaOS WIP module, comprehensive rogue detection and containment is provided.

- **Wired network management:** AWMS is capable of monitoring and managing wired switches and it can detect rogue APs and troubleshoot AP connection issues from a single, integrated platform.

Mobility Controller Deployment Options

With a solid understanding of master/local operation, it is now possible to look at the various deployment models that the system allows. The mobility controller is flexible enough to be deployed in multiple settings. In some cases the master/local functionality is collapsed into a single device for a small deployment, and in other cases the mobility controller acts as a gateway for remote devices.

Campus Deployments

Campus deployments are extremely common for Aruba solutions. The highly scalable nature of the system allows for larger deployments of multiple systems, where users can roam between floors and buildings and remain connected. Advanced services such as voice and video distribution are often integrated into these networks, which delivers an experience previously possible only with wired devices. Most deployments involve multiple local mobility controllers with redundancy deployed either in the distribution layer or data center. These deployments also have redundant master mobility controllers and AWMS in the data center.

Figure 12 *Campus deployment model*

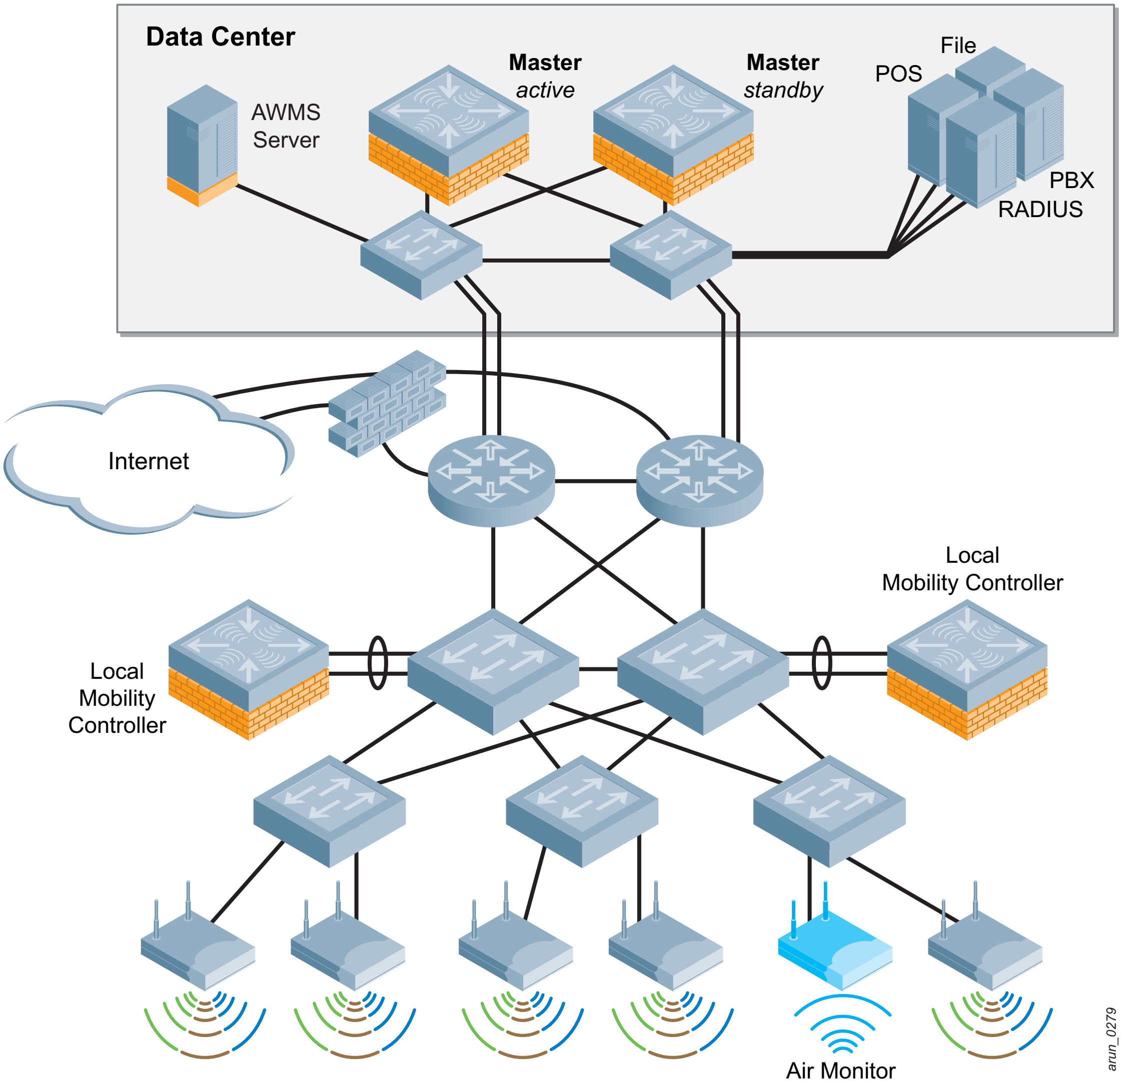

Stand-Alone Mobility Controller Operation

The stand-alone mobility controller option is a hybrid of the master and local mobility controller functionality. In this case, the mobility controller is configured as a master mobility controller, but it is also responsible for all local functionality. These devices are typically deployed in smaller organizations where only a single controller is required, and redundancy is not needed or cost effective.

Figure 13 *Stand-alone mobility controller model*

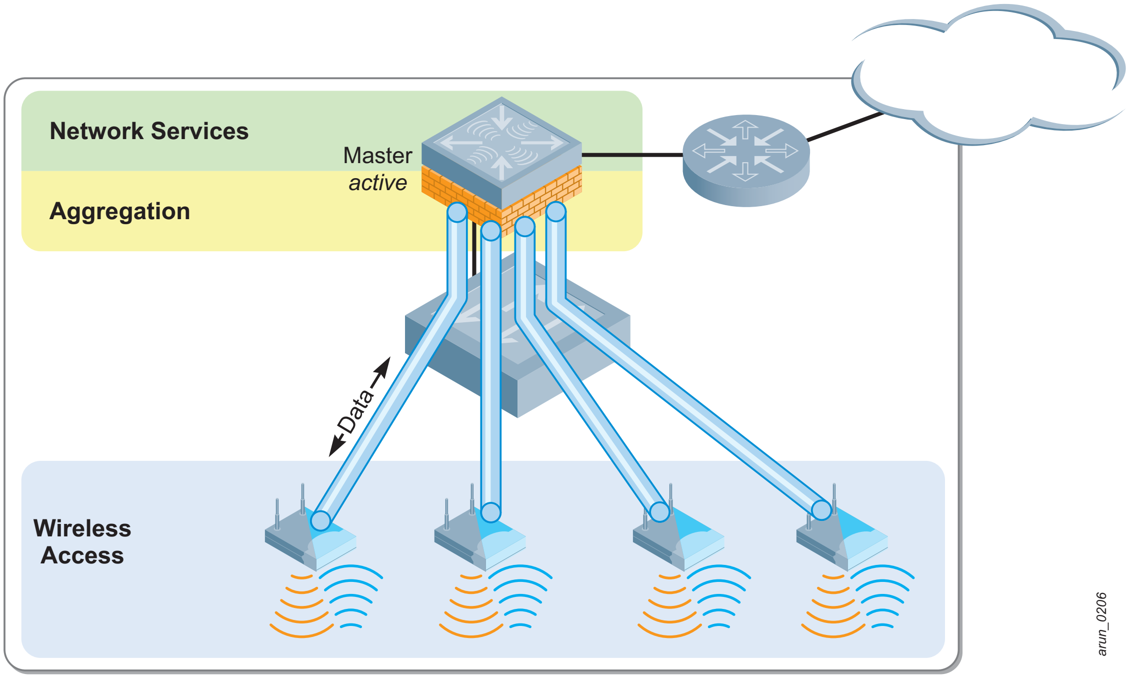

Branch Office Controller

In the distributed enterprise network, branch offices of many sizes exist. When a branch office grows beyond the capabilities of a RAP deployment, a smaller scale mobility controller that can handle multiple APs can be deployed. The Aruba product line includes scaled versions of the mobility controller family with additional branch features to fill the need of the small branch office. These mobility controllers include additional functionality for local printers, NAS devices, and 3G dial backup connections using USB modems. These devices can be deployed in stand-alone mode, as local controllers with a centralized master, or an independent master/local cluster that uses the AWMS to synchronize configurations across sites.

Figure 14 *Branch office controller model*

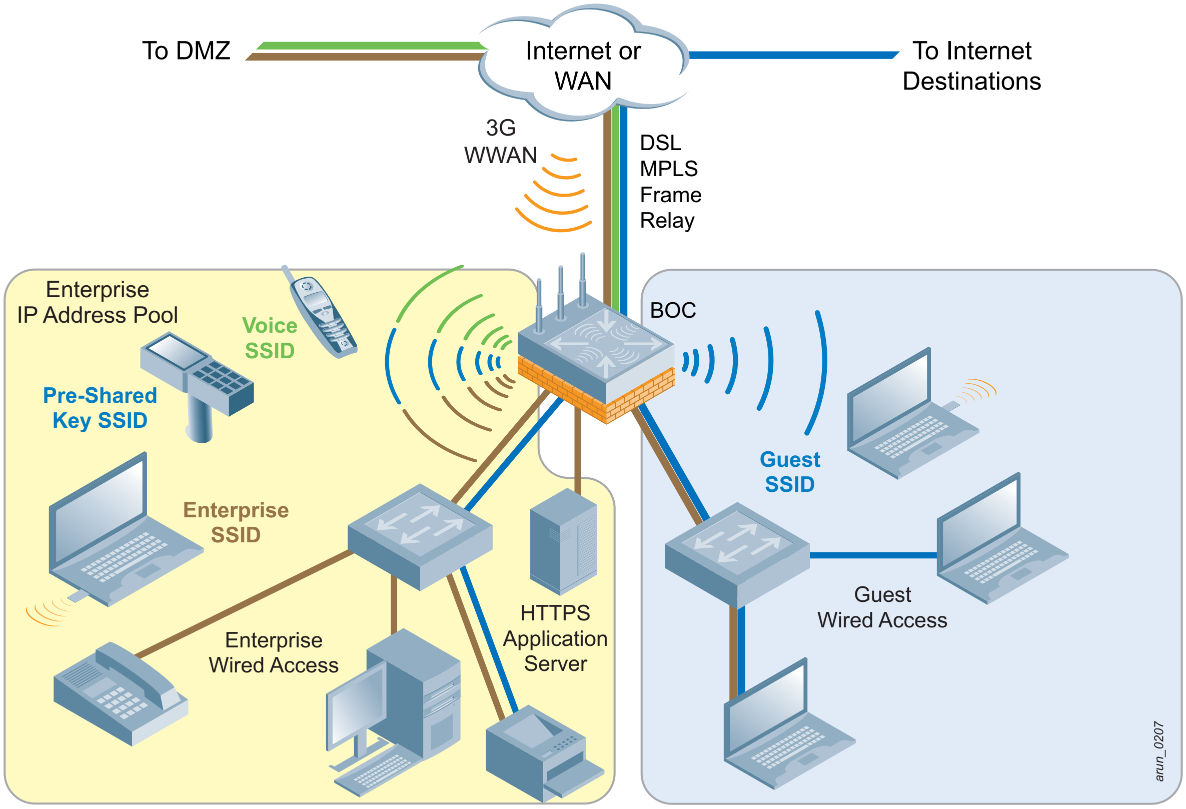

RAP and IPsec Gateway

The Aruba VIA agent, RAPs, site-to-site VPNs, and third-party IPsec clients typically terminate on local mobility controllers in the network DMZ. If this Aruba deployment is the only one in the organization, mobility controllers may be deployed in a master/local cluster in the DMZ. Smaller deployments may operate with redundancy between the local and the master, with APs and clients failing over to the master in the event of an outage. Much like how campus-based APs and AMs are terminated, these mobility controllers terminate these remote devices coming in over the Internet with IPsec-protected sessions.

Organizations may also provide guest access on an Aruba solution and want to have guest user sessions tunneled to an Aruba Mobility Controller over GRE. The RAP and IPsec gateway can fill that role. Guest user sessions are tunneled out of the internal network inside a GRE tunnel, and then routed normally from the DMZ.

Figure 15 *Remote access model*

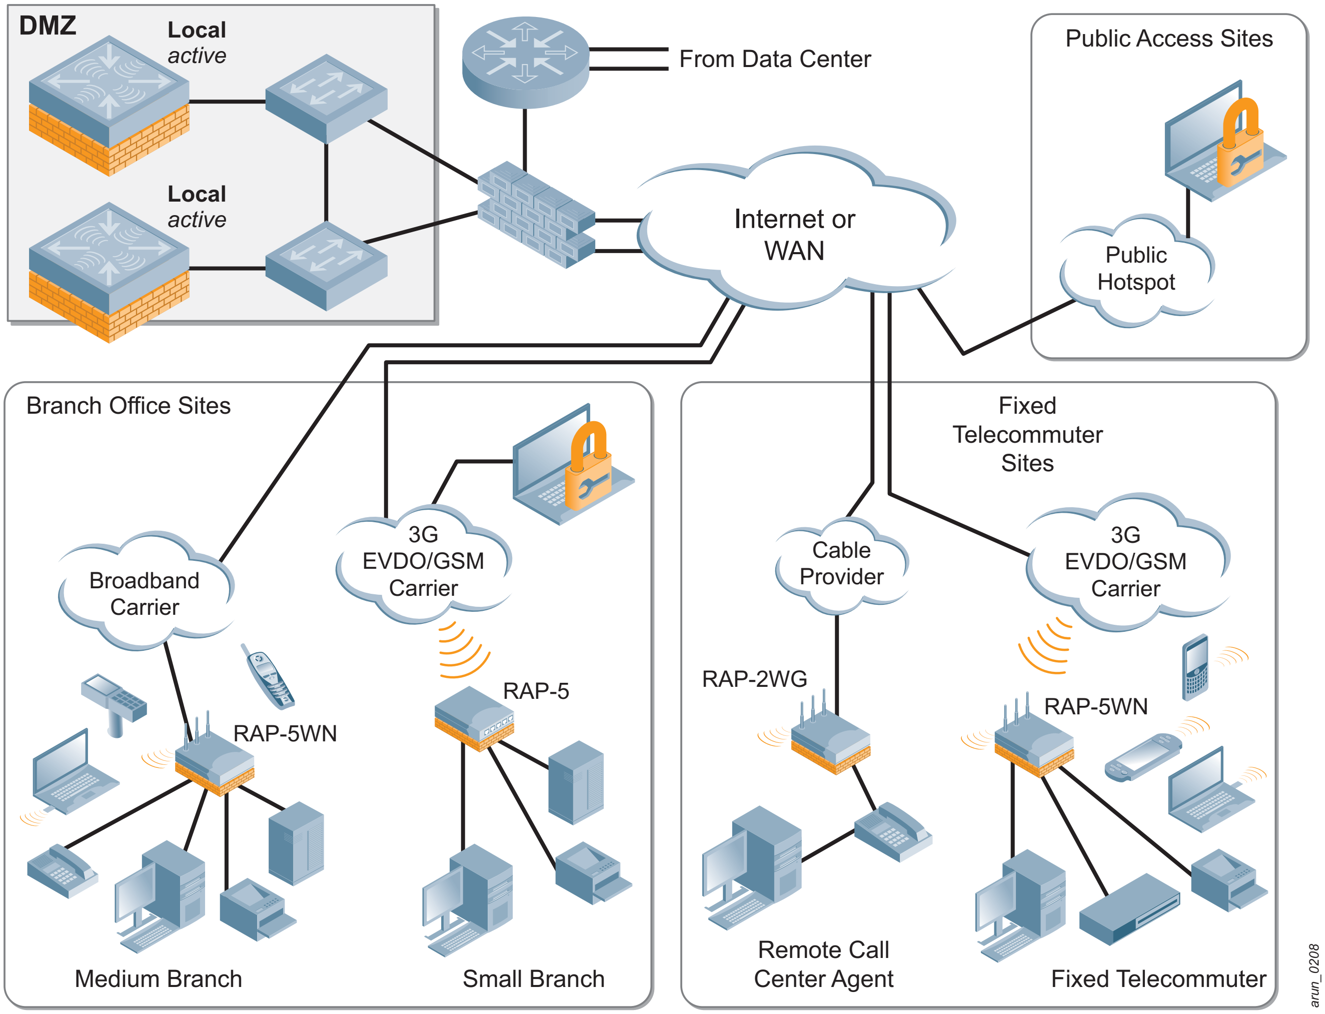

Wired Gateway

The Aruba Mobility Controllers can also act as wired gateways. The mobility controller can be a dedicated wired gateway, or it can also service remote or campus clients at the same time. By deploying a wired gateway, the organization can provide the same role-based access control (RBAC) used on the wireless network for the wired side, including secure guest access and authentication. For wired gateways, Aruba recommends that the Aruba 6000 chassis with M3 blades be used for maximum throughput. The Aruba 6000 supports up to 20 Gbps of firewall throughput per M3, and it can reach a maximum of 80 Gbps per chassis when fully populated with four M3 mobility controllers.

Figure 16 *Wired gateway model*

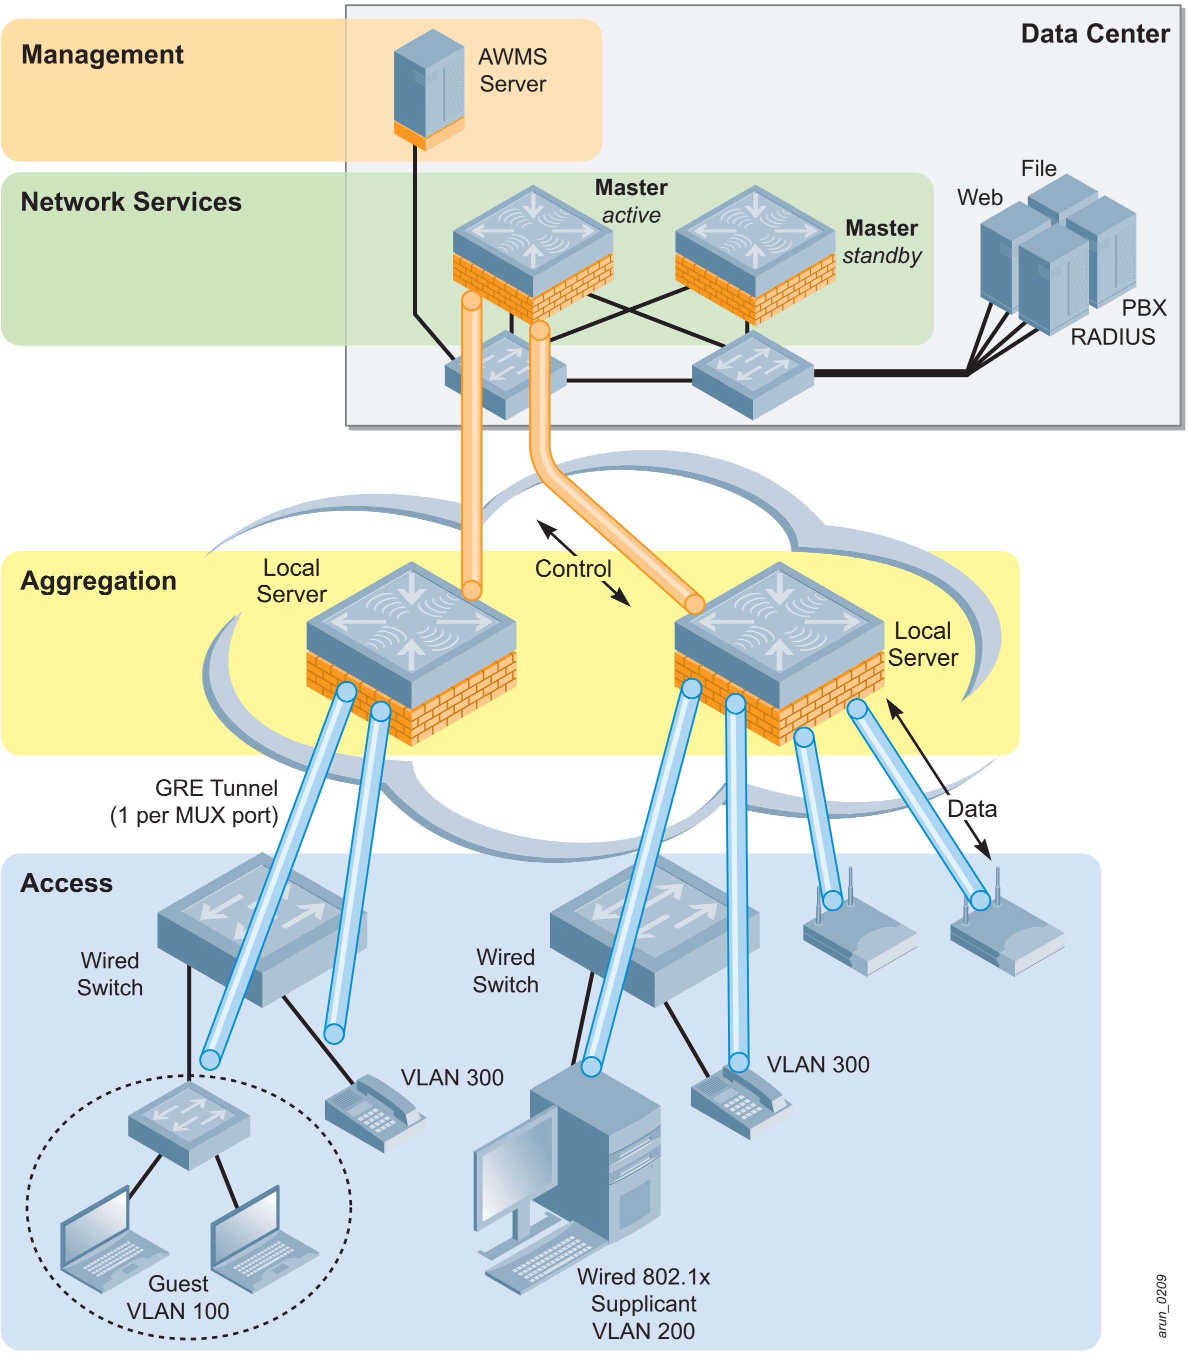

All Masters with AWMS

The Aruba concept of having master/local clusters was designed to give administrators a built-in mechanism for deploying scalable WLANs without additional tools. When AWMS is included, the network hierarchy of some deployments can be flattened by having all of the mobility controllers operate as master mobility controllers. This is not appropriate or recommended for all deployments at this time. The following sections outline how the solution is deployed.

Distributed All Masters

In distributed organizations where only a single controller will be at any one location, the preferred method of deployment is distributed all masters. This deployment essentially provides a stand-alone mobility controller that connects back to a central site. In this case, AirWave provides the consistent configuration and monitoring across the network that would have been provided by a traditional master mobility controller. This configuration simplifies the operation of the network. The site operates in a locally autonomous manner by providing local AP and device termination, but it still benefits from centralized configuration and monitoring.

Figure 17 *An all masters deployment on, separate sites*

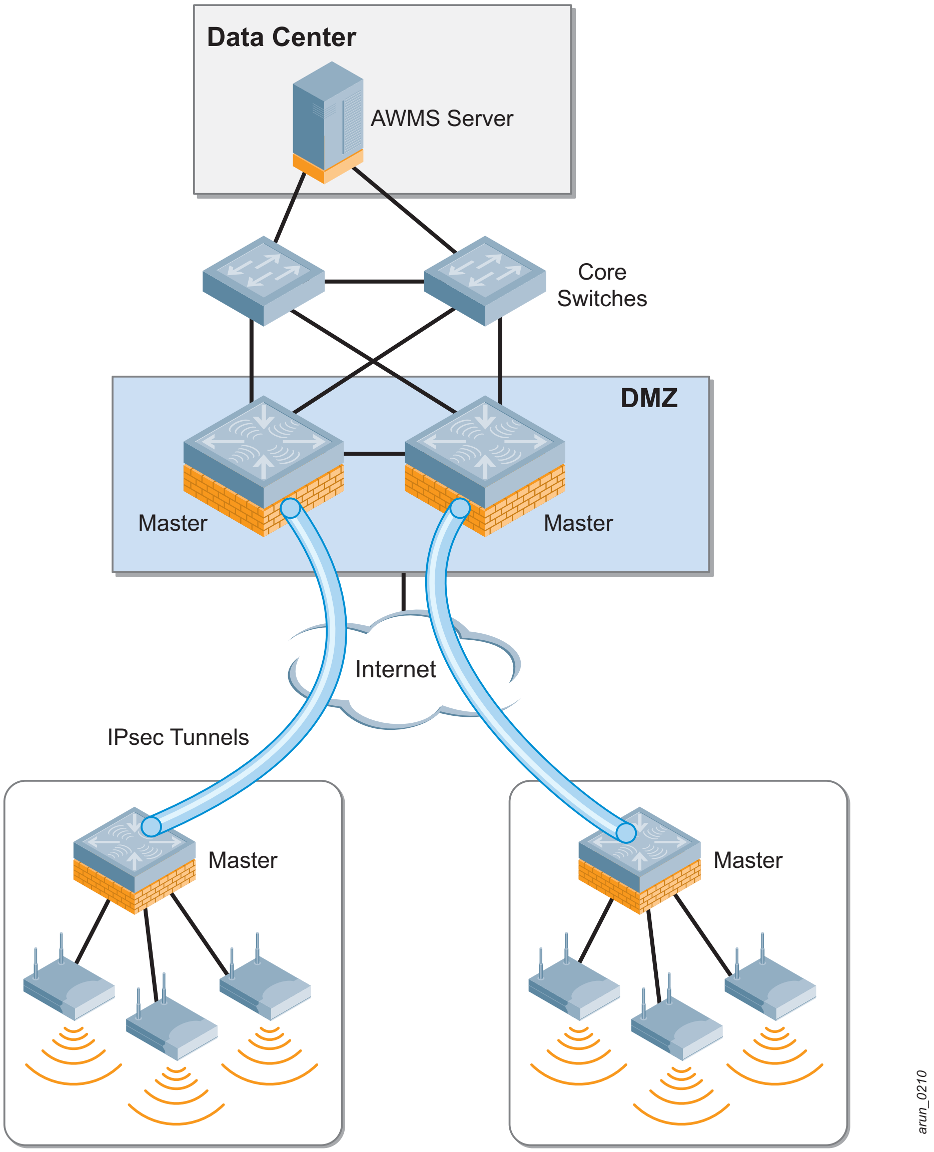

Campus All Masters

When two or more mobility controllers must coexist, additional challenges arise to the deployments if the mobility controllers are deployed as all masters. Today Aruba recommends only all master networks where the controllers are not co-located, or where all of the APs and mobility controllers in the network can be managed by a single AWMS installation. Before considering a campus-based all-master network, see Appendix C, "Scalability in the All Masters Deployment Model" on page 123, which discusses the scalability and functionality of the WLAN in a co-located all master mobility controller network based on ArubaOS 5.0 and AWMS 7.0. If these features are not required and the total device count is less than the AWMS maximum, an all master deployment can be considered.

Figure 18 *An all masters deployment, co-located controllers on a single campus*

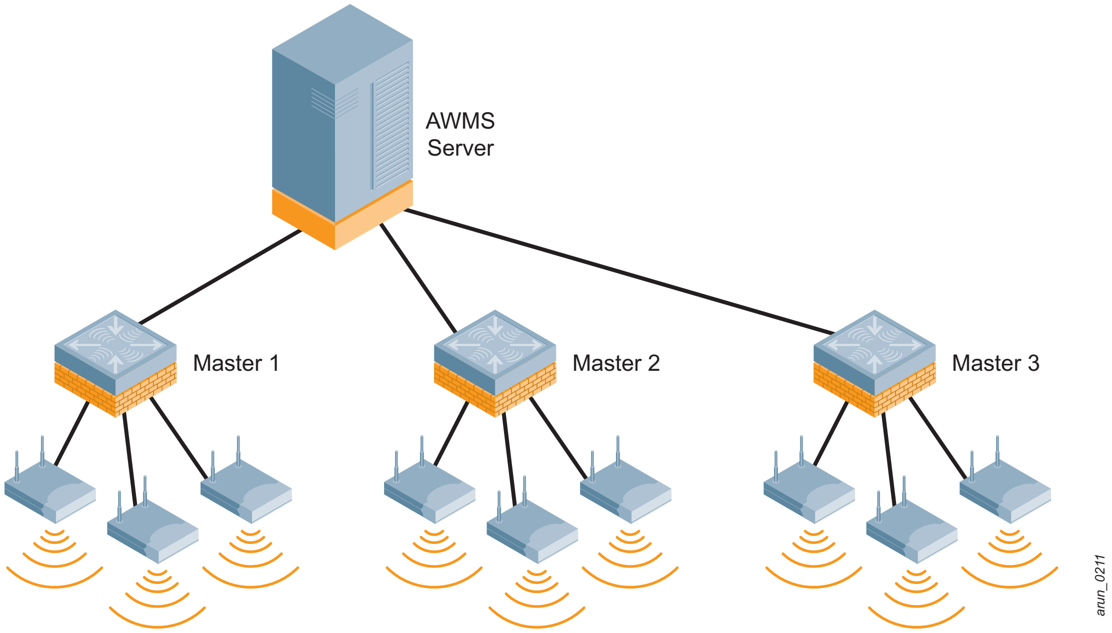

Chapter 4

Controller Licensing

The ArubaOS base operating system contains many features and extensive functionality for the enterprise WLAN network. Aruba uses a licensing mechanism to enable additional features and to enable AP capacity on controllers. By licensing functionality,   organizations are able to deploy the network and functionality in a manner that best suits their goals.

License Descriptions

- **AP Capacity:** AP capacity relates to how many APs, AMs, RAPs, and mesh points that serve clients can connect to a particular mobility controller. This license covers any AP (campus, remote, and mesh) that will broadcast a user SSID and any AMs. For mesh APs, where wireless is used for wired traffic backhaul, the mesh links that do not broadcast an SSID are not counted against this license. If the AP acts as a mesh node and an access point for users, the AP counts against the AP capacity license. Controllers can be purchased with the AP capacity already installed, and it can be upgraded later as the network grows. When planning for redundancy, the AP capacity must match the maximum number of APs that could potentially terminate on the mobility controller.
- **Policy Enforcement Firewall–Next Generation (PEF-NG):** The Aruba PEF-NG module for ArubaOS provides identity-based controls to enforce application-layer security, prioritization, traffic forwarding, and network performance policies for wired and wireless networks. Organization use PEF-NG to enforce network access policies that specify who may access the network, which areas of the network they may access, and the performance thresholds of various applications. Administrators can build a unified, integrated system for network policy enforcement by leveraging the open APIs of PEF-NG to external services such as content security appliances, network access control (NAC) policy engines, performance monitors, and authentication/authorization servers. PEF-NG is licensed by AP count, and the number of licensed APs must be equal to the AP capacity license of the mobility controller. To enable PEF-NG on wired-only gateways, a single AP PEF-NG license is required.
- **Policy Enforcement Firewall–VPN (PEFV):** The PEFV license provides the same features and functionality that PEF-NG does, but it is applied to users coming in over VPN connections as opposed to wireless users. The user role and policy are enforced on the mobility controller and thus only affects centralized traffic. Without this license, users of VPN software, either the Aruba VIA client or any third-party software, will be given what essentially amounts to an “allow all” role that cannot be modified. The PEF-VPN license is purchased as a single license that enables the functionality up to the full user capacity of the mobility controller.
- **Wireless Intrusion Protection (WIP):** The Aruba WIP module protects the network against wireless threats to network security by incorporating WIP into the network infrastructure and eliminating the need for a separate overlay system of RF sensors and security appliances. The WIP module is integrated into the WLAN system, so it provides unmatched wireless network visibility and simplicity of operation for network administrators, and thwarts malicious wireless attacks, impersonations, and unauthorized intrusions. Clients and APs are already a part of the system, so no valid AP or user list must be maintained, because the network already knows which users and devices belong there. Additionally, many of the traditional features and attacks that are reported by traditional WIDS vendors are unnecessary due to the WIP integration with the WLAN itself. WIP is licensed by AP count, and the number of licensed APs must be equal to the AP capacity license of the mobility controller.

- **Content Security Service (CSS):** Aruba CSS provides cloud-based security for branch offices and teleworkers. CSS seamlessly integrates with the RAP and BOC product families to provide high-throughput, low-latency content security with centralized reporting and management. CSS leverages data centers around the world and provides complete protection including advanced URL filtering, P2P control, anti-virus/anti-malware, botnet detection, and data loss prevention. High-speed web logs in CSS provide a flexible and powerful way to view broad trends and per-user drill-downs of Internet activity. CSS licensing is based on three components: total user count, feature bundles, and contract length (1 or 3 years). The CSS licenses are installed on the cloud-based service platform.
- **xSec (XSC):** xSec is a highly secure data link layer (Layer 2) protocol that provides a unified framework for securing all wired and wireless connections using strong encryption and authentication. xSec provides a Federal Information Processing Standard (FIPS)-compliant mechanism to provide identity-based security to government agencies and commercial entities that need to transmit extremely sensitive information over wireless networks. xSec provides greater security than other Layer 2 encryption technologies through the use of longer keys, FIPS–validated encryption algorithms (AES-CBC-256 with HMAC-SHA1), and the encryption of Layer 2 header information that includes MAC addresses. xSec was jointly developed by Aruba and Funk Software®, which is a division of Juniper Networks®. xSec is licensed on a per-user basis.

Understanding the Functionality of PEF-NG and PEFV

The functionality contained in the two different firewall licenses can be confusing. The following table highlights the features that are enabled by each of the firewall licenses as they are installed, and how they interact with one another.

Table 5 *PEF-NG and PEFV Comparison Chart*

PEF-NG License	PEFV License	Wireless Users	VIA/VPN Users	Wired/ Third-Party AP Users	Controller Port ACLs
Installed	Not Installed	Yes	**No**	Yes	Yes
Installed	Installed	Yes	Yes	Yes	Yes
Not Installed	Installed	**No**	Yes	**No**	Yes
Not Installed	Not Installed	**No**	**No**	**No**	Only MAC, EtherType, and Extended ACLs are supported

Licensing Requirements and Recommendations

Different license capacities are required for cloud services and on each mode of mobility controller, depending on its operational requirements. Each license type should be reviewed to determine if the features and functionality meet the goals of the organization, and from there the feature licensing levels can be determined. For complete descriptions of the features enabled by these licenses, visit the Aruba website at http://www.arubanetworks.com/.

Matching AP-Based Licenses

AP-based licenses should always have the same AP count when in use. Licenses that are required for backup mobility controllers must be included in this count. These licenses are AP capacity, PEF-NG, and WIP. If the mobility controller is licensed for more than one of these functions, the licensing rule is:

AP capacity = PEF-NG = WIP

As an example, if 64 AP capacity license was purchased and the organization wants to deploy PEF-NG and WIP, those licenses should be purchased to match the 64 AP capacity. The final license count would be 64 AP capacity, 64 PEF-NG, and 64 WIP.

Master Mobility Controller Licensing Requirements

The master mobility controller must manage the functionality for all other platforms, so the master must have the same license types as the local mobility controllers. Licensing unlocks configuration capabilities on the system. However, the master will not terminate APs or devices, so the master can be licensed at a much lower level than the local mobility controller where APs, AMs, and devices will be serviced.

The recommended licensing levels for masters that do not terminate users are listed in the following table. Note that only the functionality that is being enabled needs to be licensed. As an example, xSec is primarily only deployed in Federal Government and military installations, and it is not required unless it will be in use at the organization.

Table 6 *Master Mobility Controller Minimum Licensing Levels*

License	Capacity
AP Capacity	0
PEF-NG	1
PEFV	1
WIP	1
CSS	N/A
xSec	1

Local Mobility Controller Licensing Requirements

Local controllers must be licensed according to the number of devices or users that consume licenses. The license consumption table that follows describes how the different licenses are consumed on local mobility controllers that terminate user sessions and APs.

Mobility controllers should be licensed at the maximum expected capacity for that mobility controller. For instance, in a failover scenario the backup controller must be licensed to accept all the APs that it could potentially end up hosting if a failure occurs, even if that is not the normal operating level.

As an example, a pair of Aruba 3600 Series Mobility Controllers is operating as local mobility controllers. Each terminates a 40% load of APs, but acts as the backup for the APs on the other controller. Each mobility controller must be licensed to 80% of maximum capacity. If one mobility controller fails, the other must be able to add the additional APs from the failed controller. Therefore, the mobility controllers must be licensed up to the maximum capacity.

Table 7 *Local Mobility Controller Licensing Levels*

License	Capacity
AP Capacity	Any AP (campus, mesh, or remote) broadcasting an SSID, or any active AM. Mesh APs that do not broadcast an SSID (such as a point-to-point bridge) do not count against this limit.
PEF-NG	Any active AP (campus, mesh, or remote) or AM, this license must be equal to the AP capacity of the network.

Table 7 *Local Mobility Controller Licensing Levels (Continued)*

License	Capacity
PEFV	PEF-VPN is licensed by box capacity, so licenses are not consumed by individual sessions. Instead, after the license is installed, all sessions up to the box limit will have a firewall policy applied to them.
CSS	Users in the organization that have signed into the CSS service.
xSec	User sessions using xSec.

Chapter 5

AP Forwarding Modes

The Aruba dependent AP architecture originally operated in a strictly centralized design, with all traffic from users encapsulated in GRE and forwarded through GRE tunnels to the mobility controller for decryption and processing. While that model is still the preferred model for campus operation, the architecture now allows for multiple forwarding modes for user traffic. Modes differ based on the whether the AP is operating as a CAP or RAP. These modes are SSID dependent, and can be "mixed-and-matched" to suit the needs of the organization. These modes determine how user traffic is handled, including where decryption occurs and where role-based firewall policies are applied.

In all cases, the number of SSIDs should be minimized. The decision to add an SSID should be based on the need for one of three things:

- a new authentication type
- a new encryption type
- a new forwarding mode

In all cases, the organization should seek to limit the SSIDs to four or less per radio on each AP. Limiting SSIDs helps preserve airtime for data transmission by limiting the number of beacons per radio. Complexity is reduced for end users and network administrators.

The nature of forwarding modes can limit the capabilities of certain features. Please see the *ArubaOS User Guide* for a complete list of features available in each forwarding mode. The user guide is available on the Aruba support site at https://support.arubanetworks.com/. The following table lists the modes that are available depending on if the AP is deployed as a CAP or RAP.

Table 8 *AP Forwarding Modes*

Feature	Campus	Remote
Tunnel Mode	Yes	Yes
Decrypt-Tunnel Mode	Yes[a]	Yes
Bridge Mode	Yes[a]	Yes
Split-Tunnel Mode	No	Yes

[a] Requires that CPsec be enabled.

NOTE

CAPs that use bridge mode and decrypt-tunnel mode must have CPsec enabled on all APs. For CAPs deployed in the traditional tunnel mode of operation, CPsec is currently optional. This feature enables certificate-based encryption and AP white listing for CAPs to increase the security and authorization aspects of an Aruba system. CPsec is covered later in this document, and Aruba highly recommends that organizations plan their network such that the ability to turn CPsec on in the future is possible without reconfiguration. In the future it is likely that some features will depend on CPsec being operational in the network. Also note that CPsec is not required or deployable for RAPs connecting over the Internet. These APs are already protected by an IPsec session that is established using either certificates or a pre-shared key (PSK).

Tunnel Mode

Tunnel mode is traditional method Aruba uses to move traffic between the AP and the mobility controller. In this mode, a GRE tunnel is established between the AP and the mobility controller. When an AP receives a wireless frame, the AP encapsulates the frame into GRE without decrypting or modifying it. The AP sends the frame to the mobility controller. When the mobility controller receives the frame, it performs the decryption operation, applies the user's firewall policy, and forwards or filters the frame as appropriate. By centralizing encryption and decryption at the mobility controller, network security is enhanced because encryption keys are never sent to the APs. The keys are securely stored on the mobility controller.

Figure 19 *CAP tunnel mode deployment*

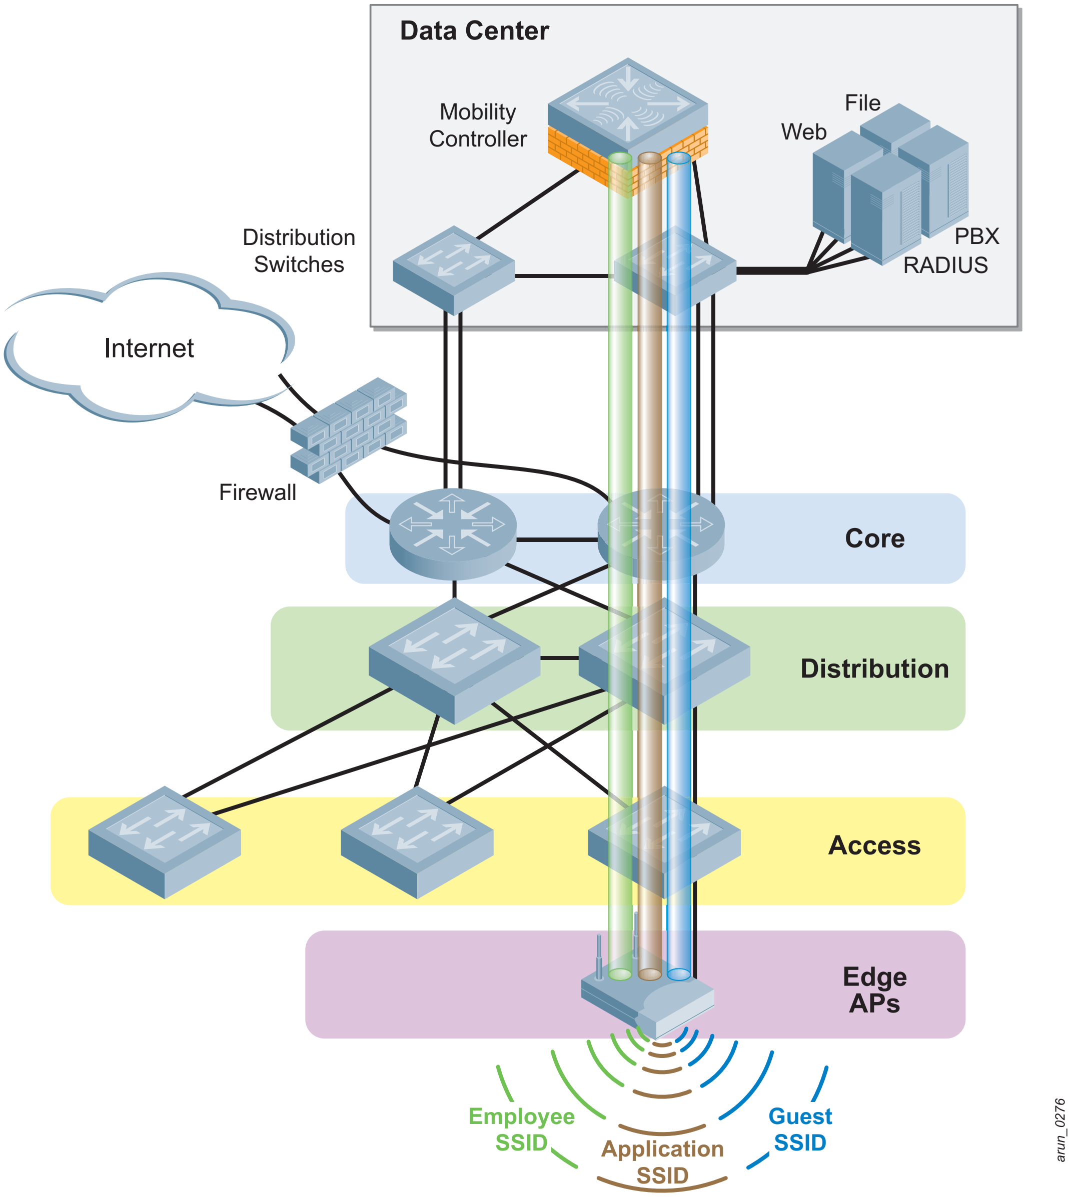

The traffic is not bridged directly onto Ethernet, so user VLANs need only be available at the mobility controller, as opposed to setting up trunked VLANs to each AP. APs are deployed directly into existing edge switches with no special VLANs required as long as the mobility controller IP is routable from the edge of the network. Tunnel mode also offers a higher level of security, because traffic is encrypted from the client to the mobility controller.

Decrypt-Tunnel Mode

Decrypt-tunnel mode is similar to tunnel mode in that all traffic transits back to the mobility controller. The difference is that the decryption of user traffic occurs on the AP before the traffic is encapsulated in GRE. This mode is primarily used to allow inline security appliances to view traffic as it flows through the network before it is filtered by the mobility controller. Users of this functionality include banking and government organizations with strict data recording mandates. Decrypt-tunnel mode can also be used for debugging by allowing traffic to be captured and inspected between the AP and the mobility controller. To enable decrypt-tunnel mode, CPsec must be enabled in the network. CPsec protects the encryption keys as they move between the controller and the AP.

Figure 20 *CAP decrypt-tunnel mode deployment*

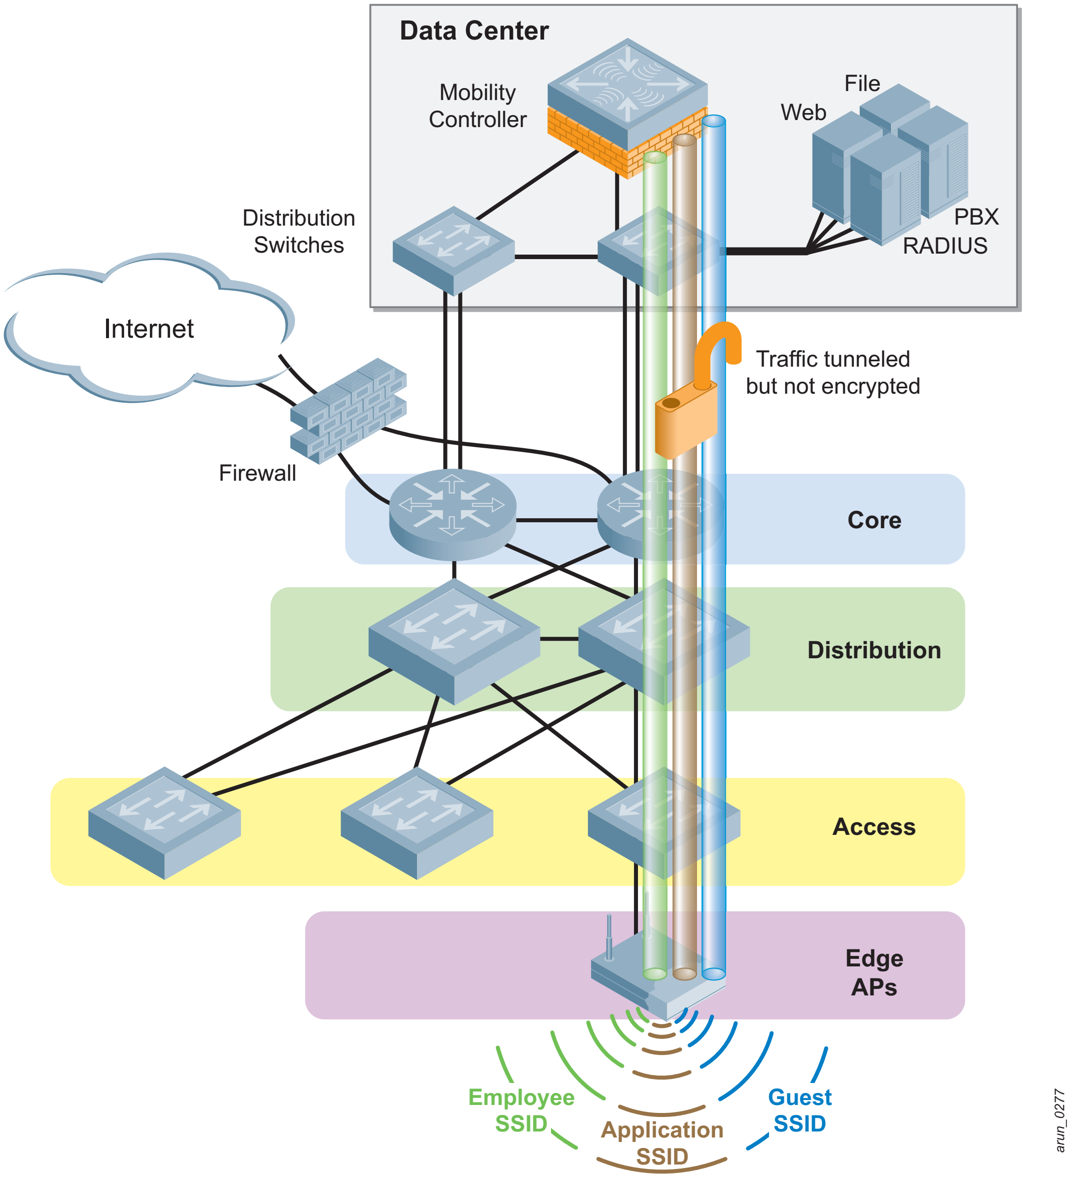

Bridge Mode

Bridge mode allows the AP to bridge traffic directly on to the LAN, with firewall policies applied at the AP. This deployment model is typically used in a deployment with a small number of users and APs on a single /24 subnet. Aruba supports no more than 32 APs at a single Layer 2 network without a controller being present and reverting to one of the other two forwarding modes.

This is not a mobility controller limitation, but a limitation in the number of devices that should reasonably be deployed in a single Layer 2 network. Most network administrators will keep Layer 2 segments limited to /24 subnets to control broadcast domain size. This limitation fits with the expected network size, providing approximately 222 station addresses, or approximately seven stations per AP. As an example, where multiple buildings exist in a small area, such as a school, if each building is a separate Layer 2 network, each building can have up to 32 APs deployed.

The APs still require access to the mobility controller to function, though the controller does not need to be in the same location as the APs. If the mobility controller is remotely located, the APs need a secure connection (VPN) between the sites with low latency. All processing is performed on the AP, so certain centralized features are not available. To enable bridge mode, CPsec must be enabled in the network.

Figure 21 *CAP bridge mode deployment*

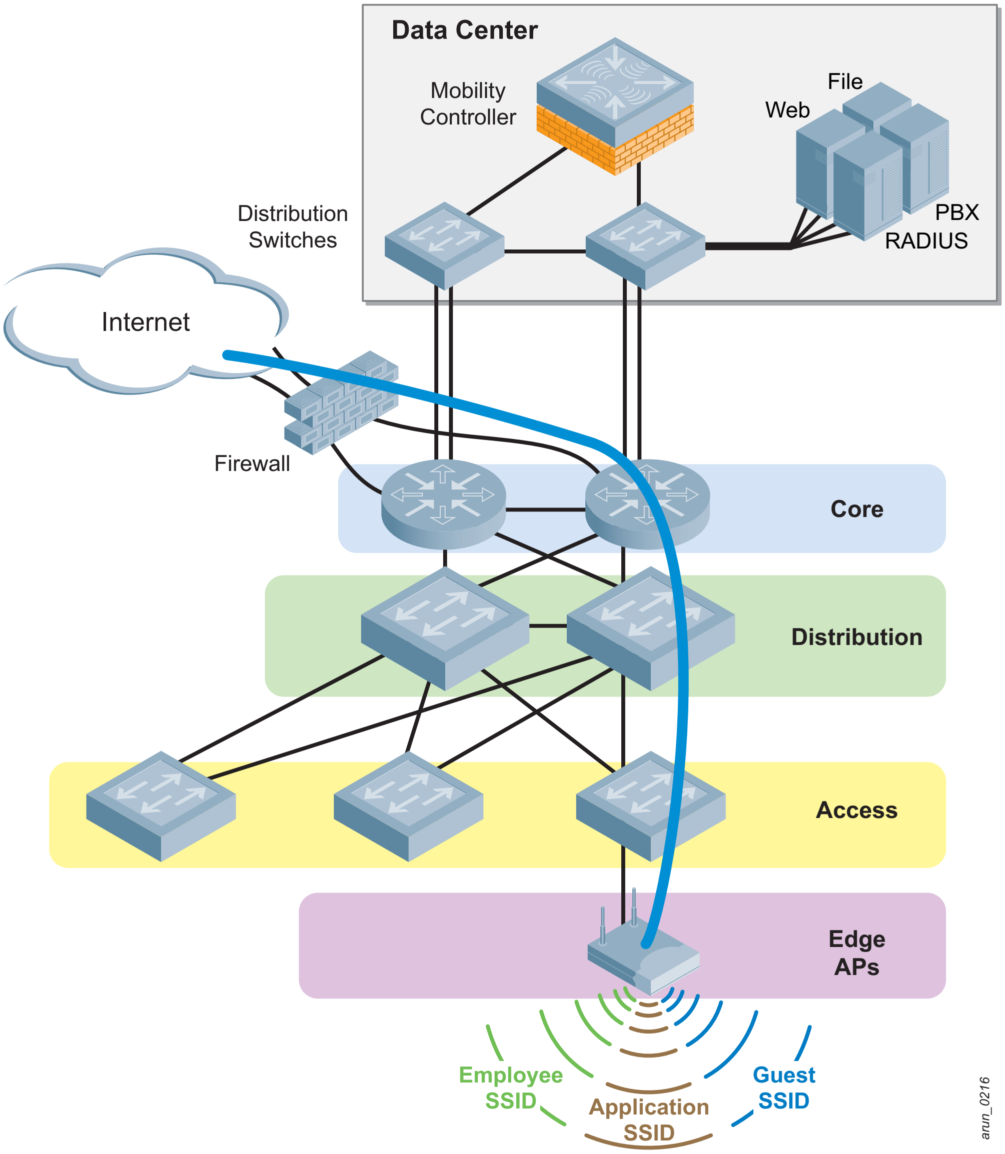

Forwarding Mode Client Limits

Encryption keys must be stored on the AP itself in decrypt-tunnel and bridge modes, so those modes support a lower number of clients per AP. Even with this reduction, the number of supported clients still far exceeds the recommended client density per AP.

Table 9 *Client Limits Per AP*

Number of Clients Per AP by Forwarding Mode and Encryption Type				
	AES	**TKIP**	**WEP**	**Open SSID**
Centralized Encryption/ Decryption (Tunnel Mode)	256	256	256	256
	Note: For centralized encryption/decryption, 256 clients is the per-AP platform limit across all radios.			
Distributed Encryption/ Decryption (Decrypt-Tunnel and Bridge Modes)	110	54	124	124
	Note: Client counts are lower for AES and TKIP due to key slot constraints.			

Campus Forwarding Mode Recommendations

The following table summarizes the Aruba recommendations for forwarding mode selections based on campus deployments.

Table 10 *Campus Deployment Forwarding Mode Recommendations*

Forwarding Mode	Aruba Recommendation	CPsec
Tunnel	Recommended deployment model for the campus.	Not required[a]
Decrypt-Tunnel	Recommended only when security requirements mandate clear text for capture or enforcement, or when troubleshooting network issues where clear text traffic between the AP and controller is considered useful.	Required
Bridge	Recommended only where a controller cannot be local to the network, but access to local network resources is still required and RAP is not a feasible option.	Required

[a] Though CPsec is not required for campus tunnel-mode deployments, the network should be designed such that CPsec could be enabled at a later date. This means that master mobility controller redundancy should be planned for. Also plan for a trust anchor if required due to network size where multiple master/local clusters are required. CPsec may be required in later revisions of code to enable new features and the network administrator should include this in the planning design.

RAP Deployments

Tunnel Mode

Tunnel mode on a RAP works in the same manner as a campus deployment, except that the RAP is operating over IPsec across the public Internet. The frame is still in its encrypted form, typically in AES-CBC in the case of 802.11i/WPA2, so IPsec places an authentication header on the packet. However, IPsec does not re-encrypt the packet, though this option is available. The packet already has been encrypted for transmission over the wireless medium, and enabling re-encryption adds overhead and processing time but yields no additional security.

Figure 22 *RAP tunnel mode deployment*

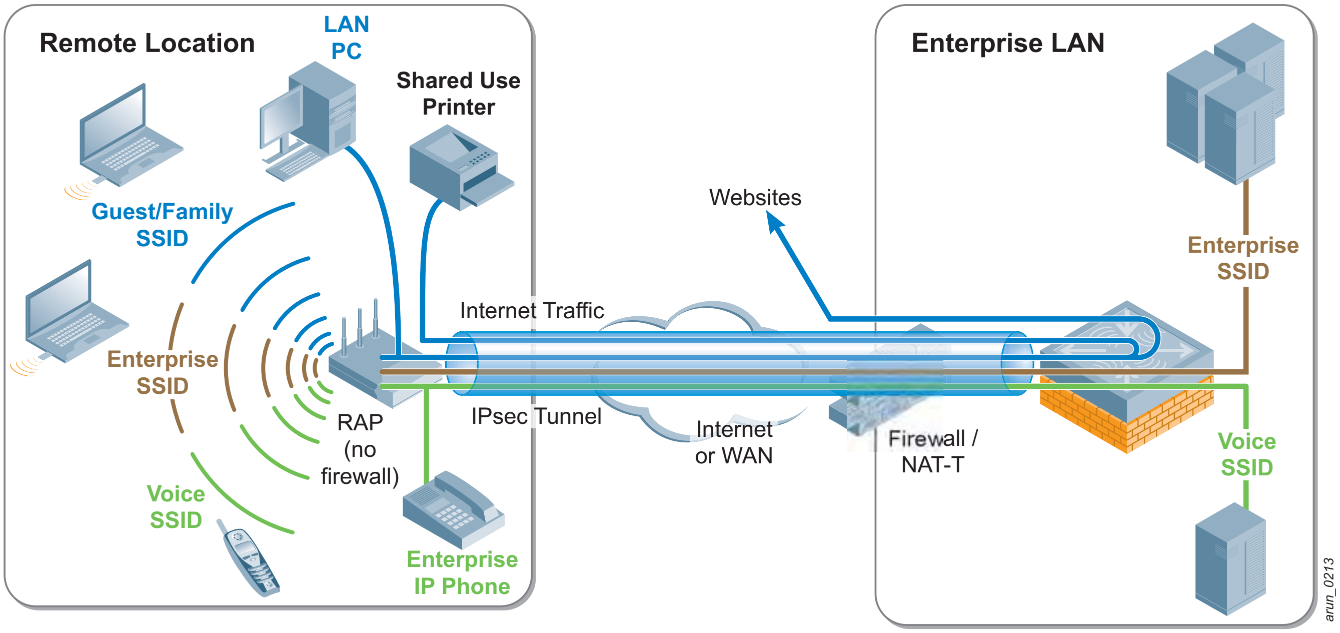

Re-encryption is recommended for any networks using weak encryption. Wired equivalent privacy (WEP) is known to be insecure and easily broken, but it is still in widespread use in many retail networks. The temporal key integrity protocol (TKIP) was introduced as a part of Wi-Fi Protected Access (WPA) as a stopgap until stronger encryption was available. However, TKIP is starting to show some vulnerability, and is on track to be removed from products certified by the Wi-Fi Alliance. Aruba recommends that use of these protocols be discontinued. While they are in use, Aruba recommends that re-encryption be enabled. Encryption must be performed for wired traffic, and the encryption is performed on the RAP.

Decrypt-Tunnel Mode

Decrypt-tunnel mode is not recommended for deployment on RAPs except where an existing VPN is in place to secure traffic. RAPs are deployed across the Internet and when decrypt-tunnel mode is enabled, the traffic is tunneled back to the controller in clear text. This mode is typically enabled on RAPs to deal with packets that arrive out of sequence due to traffic policies on the existing VPN connection between sites. An existing site-to-site VPN, such as an MPLS VPN link or dedicated VPN hardware running across the public Internet, must be in place to secure traffic between the AP and the mobility controller.

Figure 23 *RAP decrypt-tunnel mode deployment*

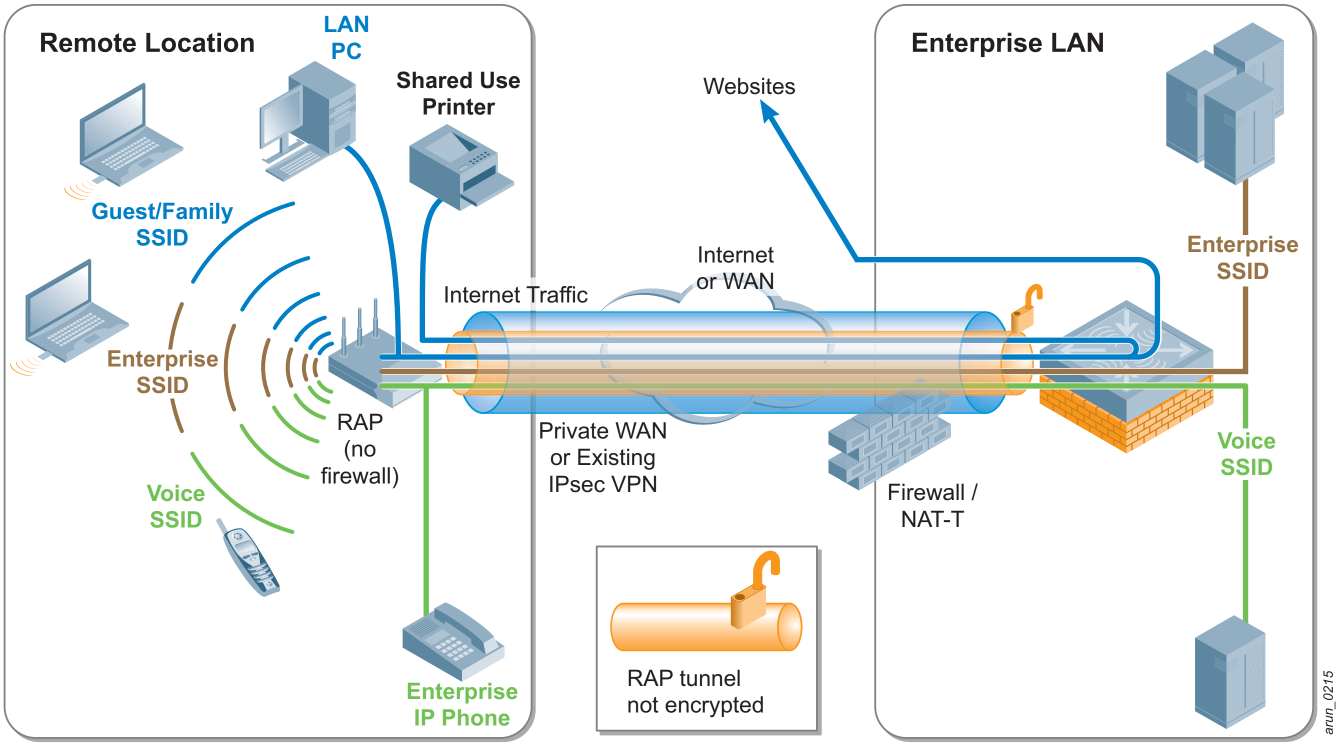

Bridge Mode

Bridge mode is available for RAP deployments, usually to provide local guest access or a connection for the user's family to leverage a single Internet connection and single AP. As with a CAP, the RAP still bridges all traffic directly to the local LAN. Even if a secure tunnel exists, users are not able to access centralized resources. Bridge mode connections on a RAP do not support captive portal authentication.

Figure 24 *RAP bridge mode deployment*

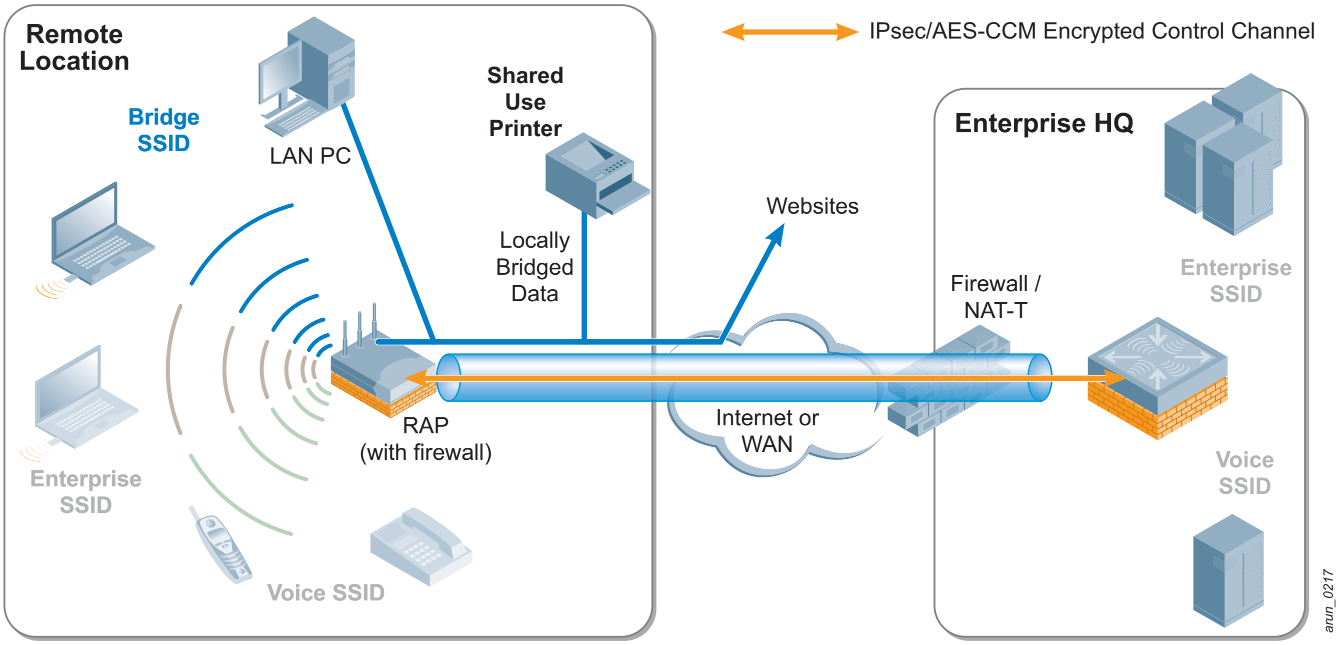

Bridge mode devices are not able to initiate communication with devices in tunnel or policy-based forwarding mode. They may initiate communication only with other bridge-mode devices, and they may not place traffic into the encrypted tunnel to the mobility controller. If policy allows, users in policy-based forwarding mode may initiate communication with bridge mode clients, such as servers or locally attached storage, and communication may occur within the bounds of the enforcement by the stateful firewall.

Policy-based Forwarding (Split-tunnel) Mode

On RAPs only, a fourth mode of forwarding is available, called policy-based forwarding. This mode has its roots in traditional split tunneling. However, instead of a simple “in the tunnel or out of the tunnel” routing decision, policy-based forwarding allows the administrator to create forwarding rules that match the internal security requirements.

Each RAP runs the same stateful firewall software that exists on the mobility controllers that determines how traffic leaves the RAP and what can enter through the RAPs Internet-facing interface. Administrators can allow certain types of traffic in the tunnel, to the local LAN, or out to the Internet.

Figure 25 *RAP policy-based forwarding (split-tunnel) mode deployment*

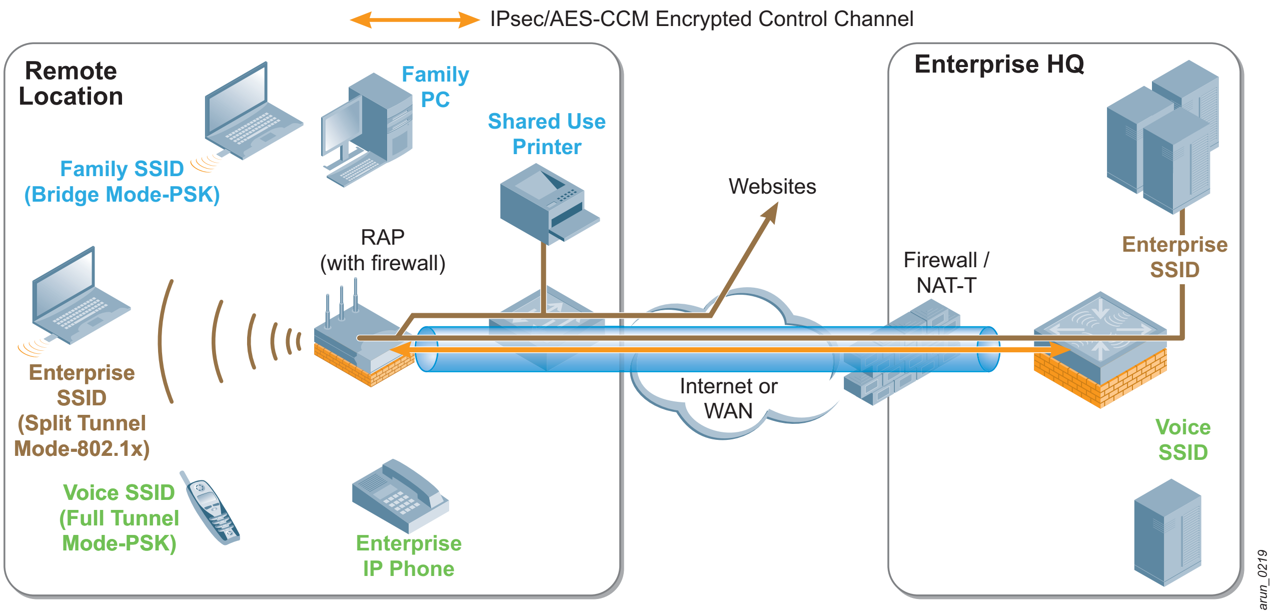

This level of granularity means that an administrator might allow printing to the local LAN segment to allow the use of a Wi-Fi enabled printer, but might block all access to the Internet directly. Traffic would be required to go through the IPsec tunnel and back to the corporate data center.

Figure 26 *Policy based-forwarding with CSS traffic inspection and enforcement*

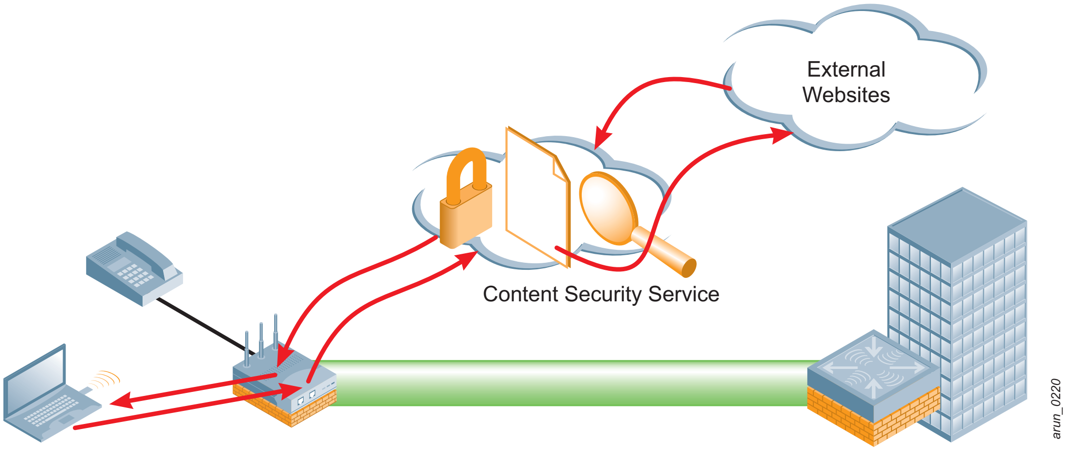

Additionally, policy-based forwarding allows for advanced cloud services such as the CSS to be enabled by identifying and acting on traffic classes to secure traffic. In this model, policy-based forwarding is used to redirect external web traffic to the Aruba CSS, which examines the request and content and enforces policy in the cloud.

NOTE

Policy-based Forwarding mode operates partially independent of the mobility controller for certain features, so some functionality is not available. This is because the APs lack visibility into the rest of the network that the mobility controller would have in a campus deployment. Please see the *ArubaOS User Guide* for a complete list of features available for Policy-based Forwarding. The user guide is available on the Aruba support site at https://support.arubanetworks.com/.

Remote Forwarding Mode Recommendations

The following table summarizes the Aruba recommendations for forwarding mode selections based on remote deployments.

Table 11 *Remote Deployment Forwarding Mode Recommendations*

Forwarding Mode	Aruba Recommendation
Tunnel	Recommended when all traffic is destined for a single central location and no local or Internet access is required unless it is first through the centralized site.[a]
Decrypt-Tunnel	Recommended only when an existing VPN infrastructure is in place to secure traffic, and all traffic is bound for the centralized mobility controller, as with the tunnel mode use case.[a]
Bridge	Recommended when providing guest access, or when a RAP is deployed in an employee's home and the family members will also use the RAP to access local resources and the Internet. No traffic crosses the IPsec connection to the central site from the bridge mode connection.
Policy-Based Forwarding	Recommended for employees. This mode provides granular control at the edge and access to local and Internet resources. This mode must be enabled when using the CSS cloud-based product. Guest access with captive portal authentication must also use this forwarding mode.

[a] This model does not allow for the use of the Aruba CSS directly from the AP. Traffic must first be forwarded to the mobility controller in the DMZ at the central site, and then it must be redirected to the Aruba CSS.

Control Plane Security

Control plane security (CPsec) is a certificate-based mechanism for protecting control plane traffic using IPsec, and optionally authorizing APs to the Aruba system. CPsec is required for campus bridge and decrypt-tunnel modes of operation. Controllers use a certificate that is issued by the master mobility controller, which uses a self-signed certificate that is unique to the local network. Each AP uses either the certificate in the Trusted Platform Module (TPM) (AP-120 Series, AP-105, AP-90 Series, and RAP-5) of the AP or has a certificate issued to the AP (AP-60/AP-61, AP-65/AP-65-wb, AP-70, and AP-85).

For very large networks with more than one master mobility controller, the networks should be linked together so that the network has a common trust anchor. The trust anchor makes white list management and AP failover between devices more simple. To enable this feature, one master should be selected to act as the trust anchor for the rest of the network. The other masters connect to the trust anchor and download a copy of the root certificate.

When CPsec is run in the network, Aruba strongly recommends that master redundancy be enabled on the master mobility controller, with a standby master mobility controller deployed. SeeChapter 10, “Redundancy Models ” on page 91 for specific recommendations for CPsec.

IPsec Connections and Certificates

When CPsec is enabled, encryption between the AP and the controller is always on. Each AP uses its built-in or controller-issued certificate to authenticate itself to the mobility controller and vice versa. When authentication is complete, the two sides construct an IPsec tunnel to protect the control channel data between the AP and the mobility controller. The traffic that is placed in the encrypted channel includes AP configuration, user encryption keys, and AP firmware updates. User traffic is not placed in this tunnel.

Figure 27 *Process for establishing CPsec between mobility controllers and APs*

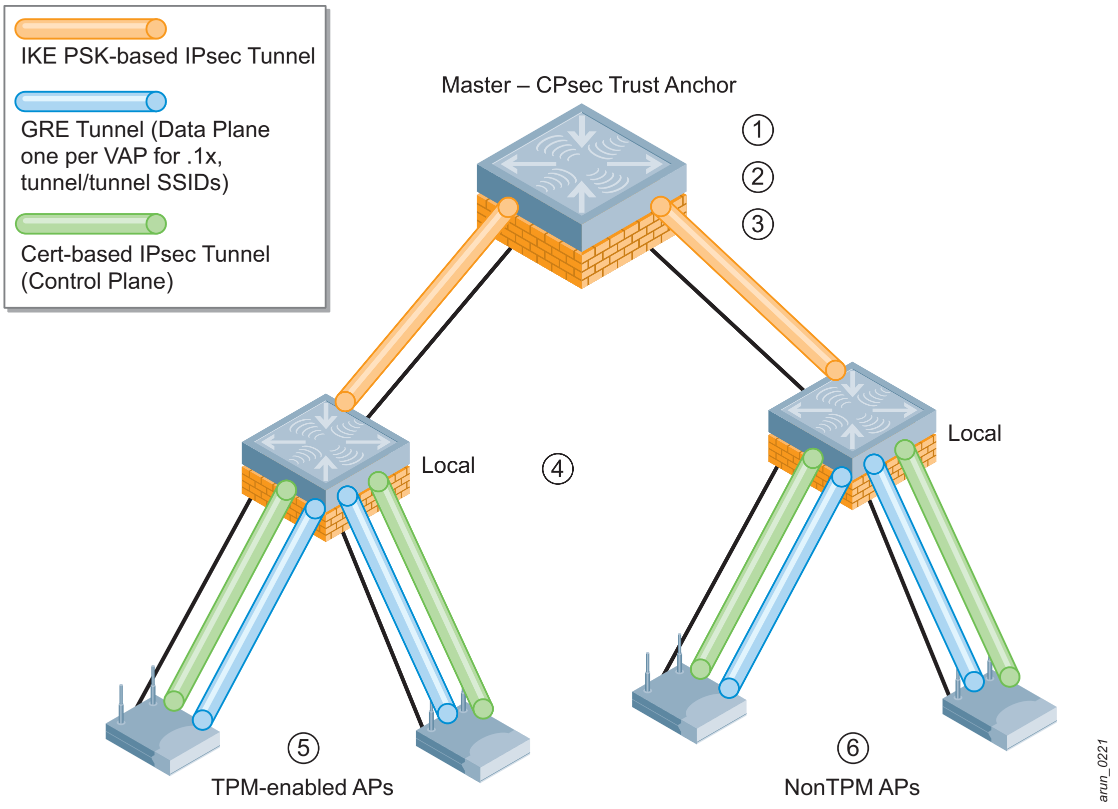
IKE PSK-based IPsec Tunnel
GRE Tunnel (Data Plane one per VAP for .1x, tunnel/tunnel SSIDs)
Cert-based IPsec Tunnel (Control Plane)
Master – CPsec Trust Anchor
1
2
3
Local
4
Local
5
TPM-enabled APs
6
NonTPM APs
arun_0221

When CPsec is enabled, the following steps occur:

1. CPsec is configured on the master mobility controller.
2. Save configuration / write memory occurs, which causes the configuration to be pushed to the local mobility controllers.
3. The local mobility controllers stop processing AP traffic and request certificates from the master mobility controller over a PSK-based connection to the master mobility controller.
4. When the local mobility controller has its certificate, it downloads the white list of authorized APs (if configured).
5. The local mobility controller begins processing requests for certificates from APs that have a certificate and are in the white list maintained by the master. The master is responsible for synchronization of all APs added to the white list from any mobility controller in the cluster.

6. If APs in the network do not have certificates, the local mobility controller issues certificates to those APs. The APs receive their certificates, establish a link to the local mobility controller, and begin broadcasting SSIDs and accepting clients.

NOTE

> When CPsec is first enabled, it can take some time to authorize all of the APs and issue all the required certificates. Aruba recommends that the organization schedule an outage window in which this provisioning can occur. Expect at least 30 minutes of downtime before all APs are certified and recovered. Schedule a maintenance window before starting a CPsec rollout in the organization.

The following tables summarize the AP types and certificates that those models of APs will use in a CPsec environment.

Table 12 *Certificate Table for 802.11n APs*

AP Type	Certificate Type	Key Length	Notes
AP-92/ AP-93	Factory installed, key stored in TPM	2048	20-year validity period starts from the date of manufacture.
AP-105	Factory installed, key stored in TPM	2048	20-year validity period starts from the date of manufacture.
AP-120 Series	Factory installed, key stored in TPM	2048	20-year validity period starts from the date of manufacture.
RAP-5/ RAP-5WN	Factory installed, key stored in TPM	2048	20-year validity period starts from the date of manufacture. **This AP is NOT supported in campus deployments.**

Table 13 *Certificate Table for 802.11a/b/g APs*

AP Type	Certificate Type	Key Length	Notes
AP-60	Issued and certified by master mobility controller	2048	20-year validity period starts from the master mobility controller start time.
AP-61	Issued and certified by master mobility controller	2048	20-year validity period starts from the master mobility controller start time.
AP-65/ AP-65-WB	Issued and certified by master mobility controller	2048	20-year validity period starts from the master mobility controller start time.
AP-70	Issued and certified by master mobility controller	2048	20-year validity period starts from the master mobility controller start time.
AP-80	Issued and certified by master mobility controller	2048	20-year validity period starts from the master mobility controller start time.
AP-85	Issued certified by master mobility controller	2048	20-year validity period starts from the master mobility controller start time.
RAP-2WG	Factory installed, key stored in flash	2048	20-year validity period starts from the date of manufacture. **This AP is NOT supported in campus deployments.**

Campus AP White List

The campus access point (CAP) white list feature enables the organization to limit which CAPs are allowed to connect to controllers in the network. By default, when any Aruba CAP is plugged into a network, it attempts to establish a connection to the master mobility controller. Many organizations want to disable this capability and instead create a white list of accepted APs for their network.

The CAP white list is always on when CPsec is enabled, but it can be configured to allow any Aruba CAP to attach to the network. Allowing any Aruba CAP to attach is the recommended operating mode for CAPs unless the organization has a security policy that mandates use of the white list. Using a white list places more burden on administrators deploying the network. If the organization wants to control the CAPs that can attach to the network, then auto-certification should be disabled following initial CAP provisioning. Auto-certification simplifies CAP deployments and allows the system to continue to operate in the traditional Aruba deployment model.

The following table summarizes the features and default state of the CAP white list feature.

Table 14 *Control Plane Security Features and States*

Feature	Initial State and Description
Control Plane Security	Disabled by default.
Auto Cert Provisioning	Disabled by default. If this feature is enabled, the controller sends certificates to all associate CAPs.
Addresses Allowed for AutoCert – All	Enabled by default. All the CAPs that are associated receive automatic certificate provisioning. CPsec must be enabled for this option to take effect.
Addresses Allowed for AutoCert – By Subnet	Disabled by default. Send certificates to a specified group of CAPs within the specified IP range.

When auto cert provisioning is enabled, the local mobility controller adds new APs to the AP white list and propagates these APs to the master's central list. This central list is then redistributed to the other mobility controllers in the network.

AP and Controller Tunnels

When the AP has its certificate, it attempts to establish a connection to the local mobility controller. The local controller checks the white list to ensure that the AP is authorized, and then presents its certificate to the AP. Both sides trust the certificates because the same authority signs them, and they use the certificates to establish an IPsec connection between the devices.

Figure 28 *CPsec tunnels between the mobility controller, AM, AP, and mesh points*

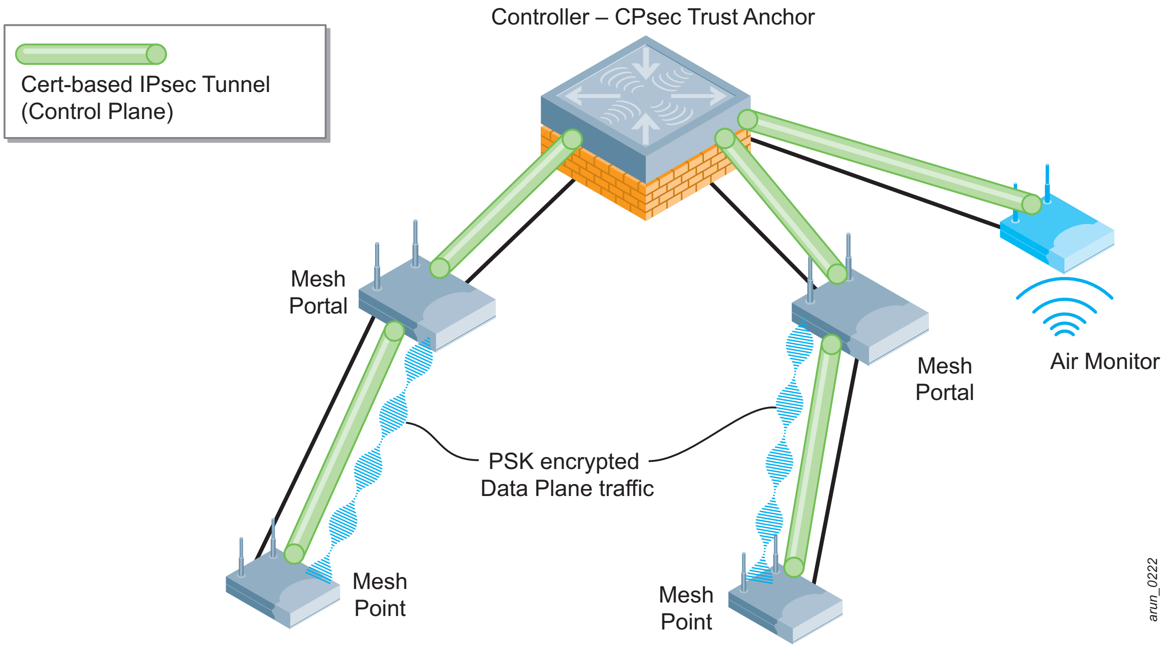

When IPsec is established, the AP receives its configuration. A GRE tunnel is established for user authentication, traffic, and monitoring to the local mobility controller for each SSID on the CAP. At this point the AP is operational, and the local mobility controller has at least two tunnels established for each AP.

Trust Anchors

When the CPsec-enabled network expands beyond a single master/local cluster, a single master mobility controller must act as the trust anchor for the entire network. The trust anchor provides the ability to have a single white list and a single set of certificates, which makes the network more flexible in terms of failover of components.

Figure 29 *CPsec on stand-alone master, master/local cluster, and multiple clusters without a common trust anchor*

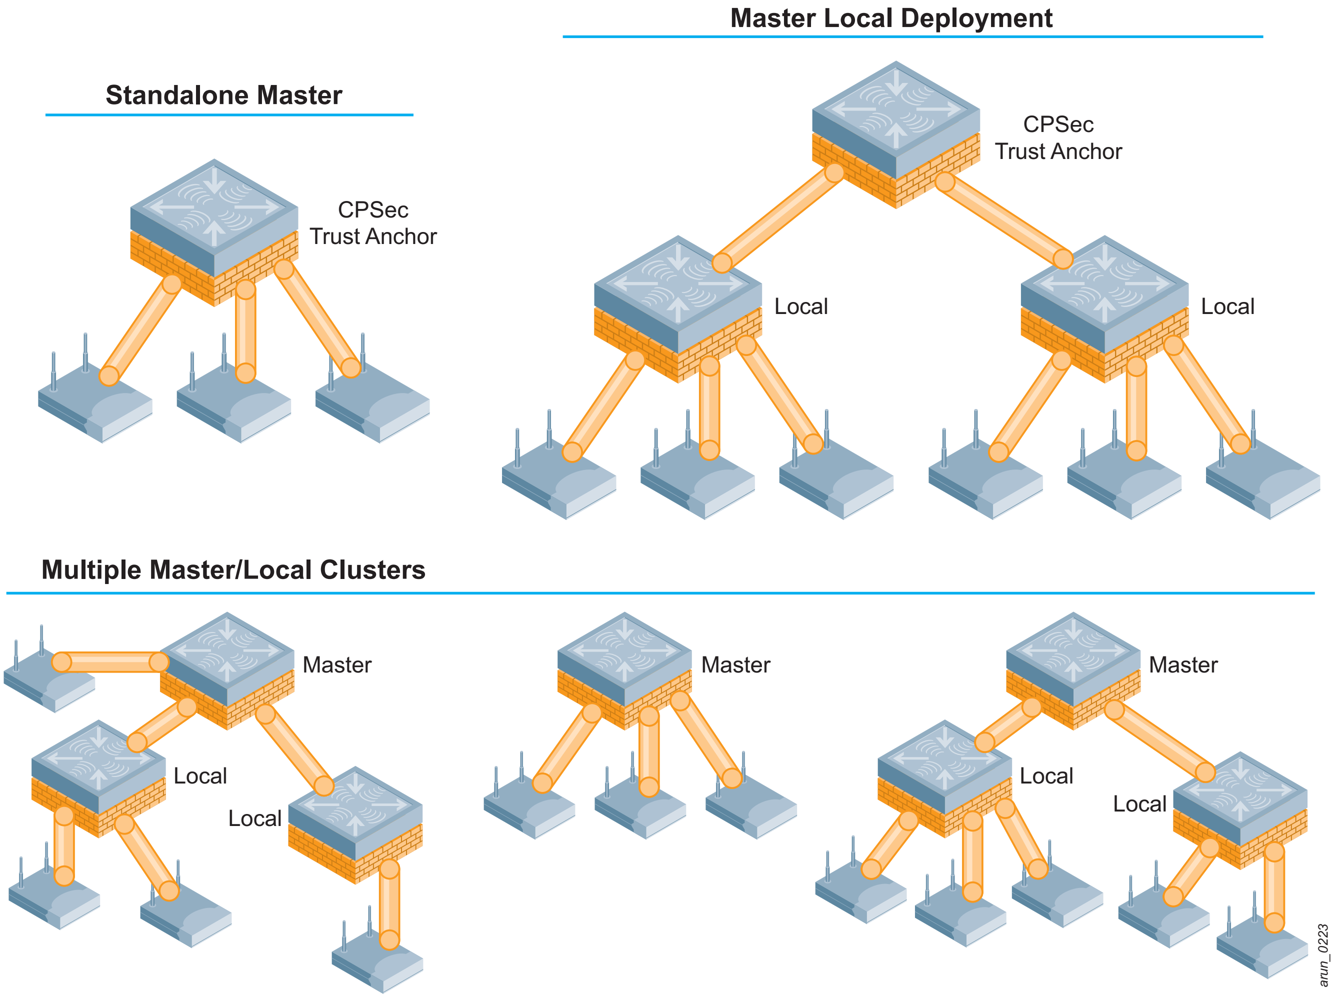

This illustration shows a series of master mobility controllers operating as independent CPsec clusters. In this mode of operation, if an AP fails between clusters, it is added to that cluster's white list and a new certificate is provisioned from the new master mobility controller.

Figure 30 *CPsec between multiple master/local clusters with a common trust anchor*

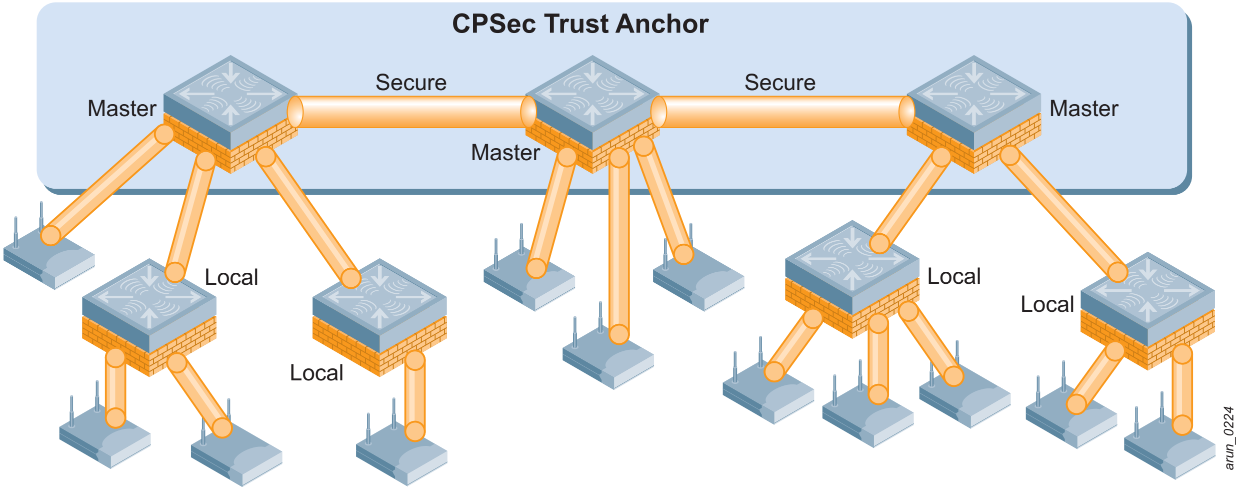

This illustration shows that all the master mobility controllers use a single trust anchor to deploy certificates to the rest of the network. In this case, all mobility controllers on the network share a common white list and one mobility controller issues certificates. All components of the network share a common certificate root and a common AP white list.

Considerations in CPsec Deployments

The following table lists the pros and cons of deploying CPsec in the organization, independent of the forwarding mode.

Table 15 *CPsec Deployment Considerations*

Feature	On	Off
Forwarding mode	All modes are available.	Only tunnel mode is available.
Control traffic encryption	Yes	No
Certificates required for AP operation	Yes	No
Boot time	Some additional boot time required, see Appendix B, "AP Failover Times" on page 121 for more information	No additional boot time required.
Permanent master failure	After the new master is back online and restored, all APs must recertify, approximately a 10-30 minute outage depending on AP type. (See Appendix A, "CPsec Scalability" on page 119 for times.)	Master must be restored, all APs continue to function.

Chapter 6

Physical Deployment Considerations

When discussing the deployment of physical hardware, Aruba uses a traditional multitier model to define the possible locations where equipment can be placed. The model consists of an access layer, distribution layer, core, data center, and DMZ.

Figure 31 *Physical deployment overview of mobility controllers*

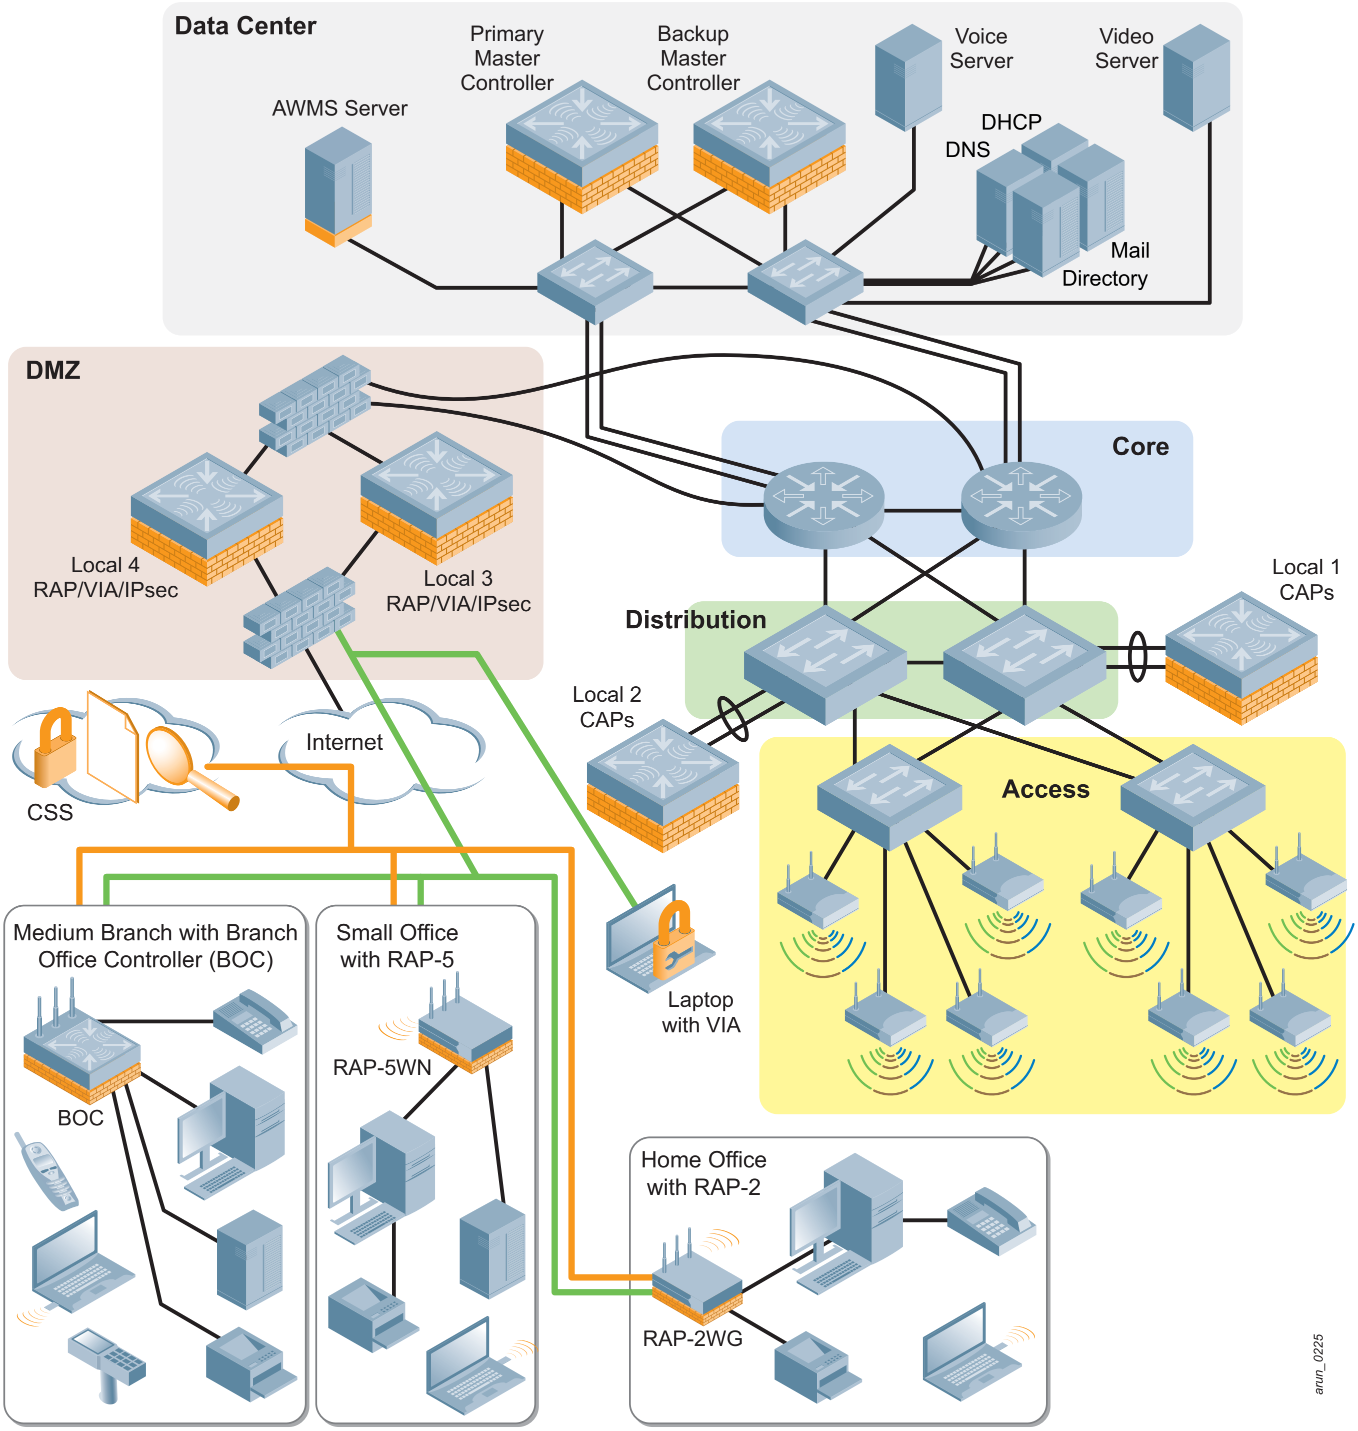

APs are always deployed at the access layer of the network, but there are more choices to be made when it comes to deploying the other components of the system.

Mobility Controller Deployments

Mobility controllers can be located in many different places, depending on the network and traffic patterns. This section explains which controllers are commonly located in each location.

Data Center

The most common Aruba component deployed in the network data center is the master mobility controller. These controllers are commonly outside the user data path, and they provide control plane functionality and services for the network. These controllers do not need to be in the data path to perform their functions, so it is best to deploy them in the data center.

Figure 32 *Redundant masters deployed in the data center with redundant links between switches*

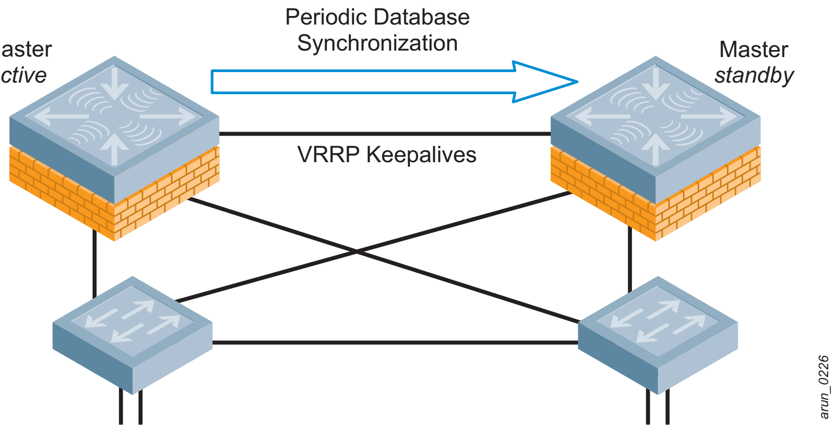

It is also common to deploy local mobility controllers in the data center in two cases:

- When the connection to the Internet is through the data center network
- When the majority of traffic is destined for the data center and not the Internet

Figure 33 *Redundant masters and locals in the data center*

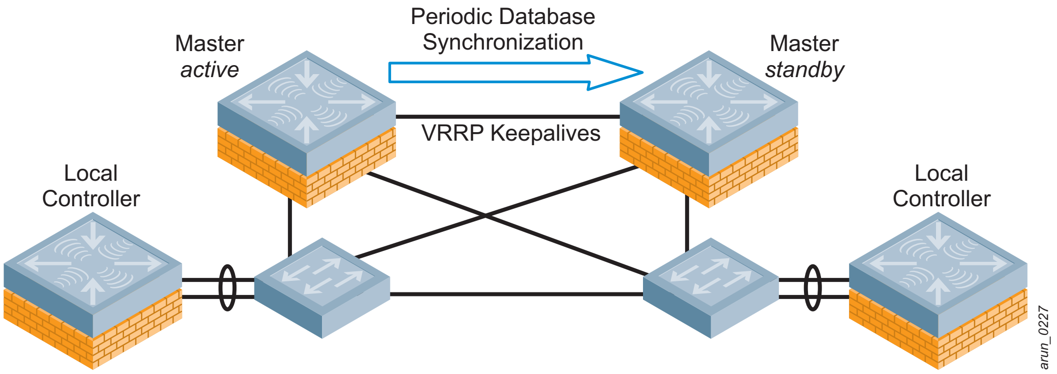

Core

Mobility controllers are typically not located on core routers, because interfaces are unavailable and network administrators want to keep the core layer as clean as possible. If the organization wants to place the controllers here, such as to move them as close as possible to an Internet connection off the core, only local mobility controllers are recommended.

Figure 34 *Redundant locals deployed at the network core*

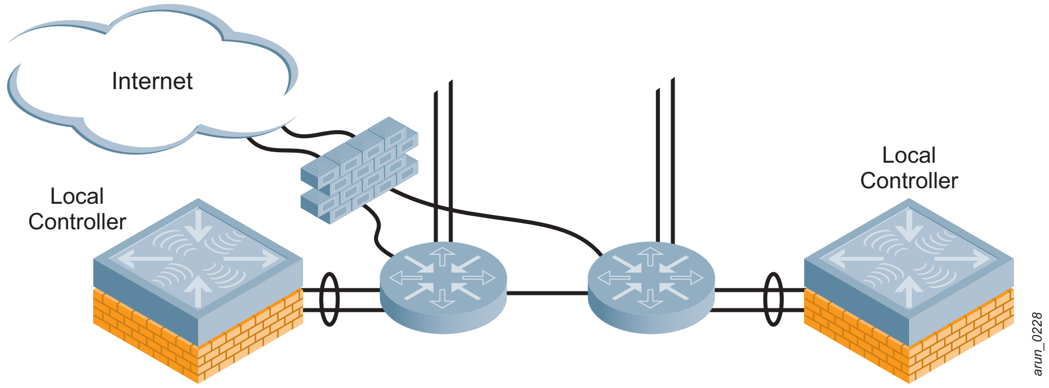

Distribution Layer

The distribution layer is the most common layer in which to deploy local mobility controllers. Typically, the distribution layer is close enough to the APs that traffic is not being switched deep into the core. This layer also gives an appropriate offload point for Internet and data center traffic, assuming that Internet traffic does not need to transit the data center network.

Figure 35 *Redundant locals deployed at the distribution layer*

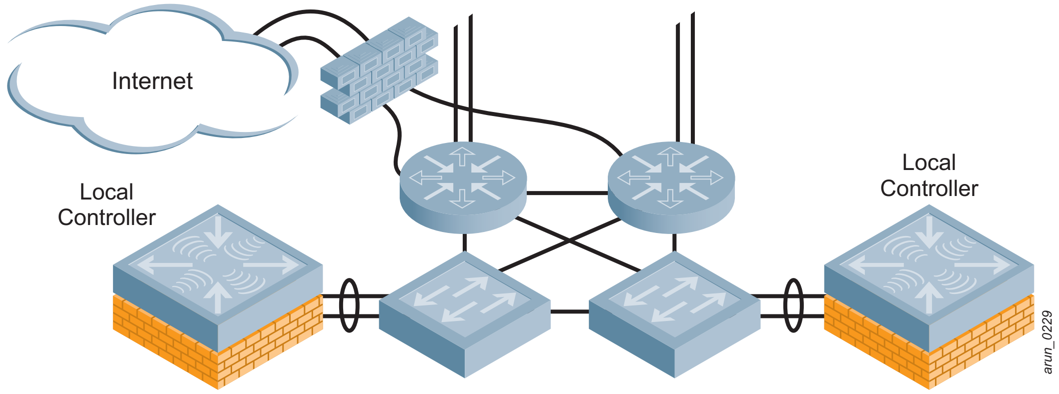

Access Layer

The access layer is the point at which end user devices directly connect to the network. In the majority of deployments, only Aruba APs, AMs, and wired APs are plugged in at the access layer. Aruba does not recommend deploying mobility controllers here unless they are providing direct, wired access and there is a requirement that users not be able to communicate with one another until they have passed through the mobility controller firewall. Such a deployment could include conference rooms where both guests and employees will use the wired network. In the majority of cases it is more cost effective to deploy a traditional wired switch and then trunk users back to the mobility controller.

Figure 36 *Mobility controllers deployed at the access layer*

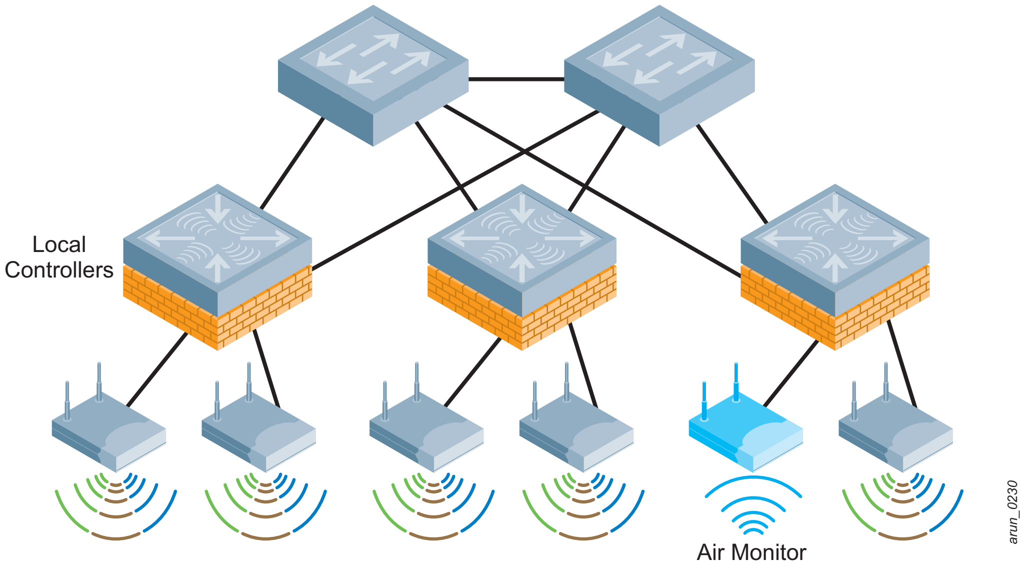

The access layer is always the appropriate choice in these two situations:

- The mobility controllers are in small, stand-alone offices where power-over-Ethernet (PoE) from the mobility controller is being provided.
- The organization is small enough that the mobility controller is the only LAN device on the site.

Figure 37 *Stand-alone mobility controller at the access layer*

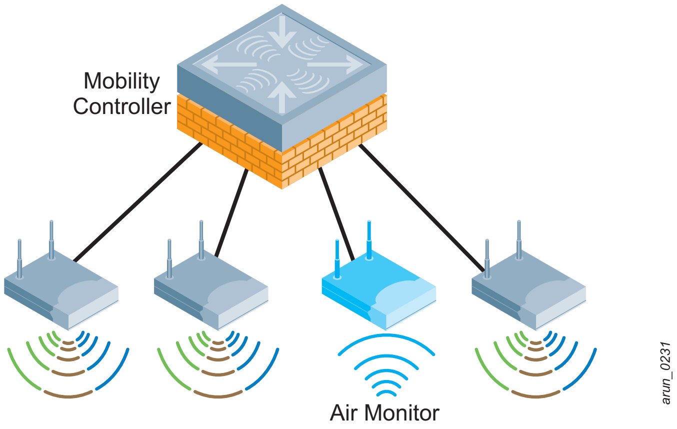

DMZ

The DMZ typically houses mobility controllers for terminating RAPs, client sessions from the VIA agent, VPN connections from third-party software, and site-to-site VPN connections from other controllers or networking devices. In most organizations, this deployment consists of one or more pairs of redundant local mobility controllers and the master controllers in the data center.

Figure 38 *Local mobility controllers in the DMZ for remote access*

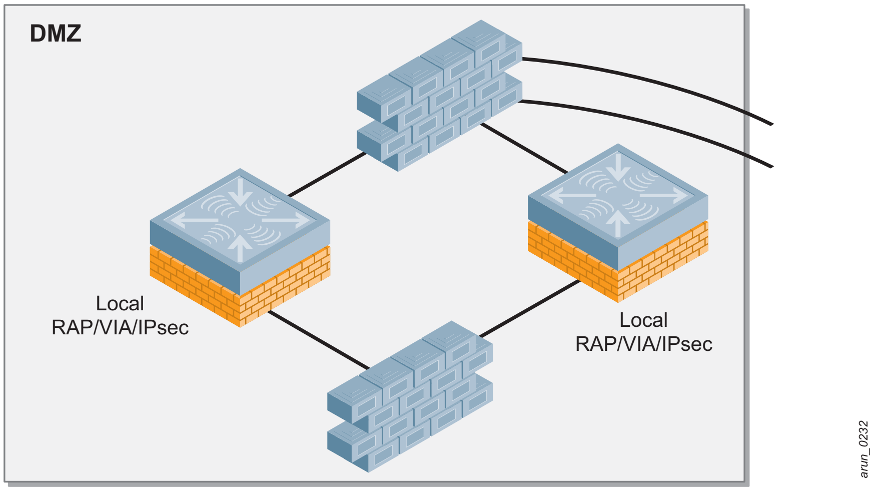

In some cases, the local and master mobility controllers are located in the DMZ. This configuration is typical of an organization that is just running Aruba as their remote access solution, or where the organization has a large enough remote and campus population to require a separate master/local cluster in the DMZ. This configuration is also common in large deployments for remote access, where the organization may have different change control windows or procedures for remote networks vs. the campus network. Deploying masters in the DMZ strictly for remote networking allows the organization to run different versions of code on their remote networks and campus, and upgrade them separately. These clusters can either be created with a separate set of masters to control multiple locals, or with only two controllers, one as a local and one as a master that backs up the local.

Figure 39 *Redundant master/local cluster in the DMZ for remote access*

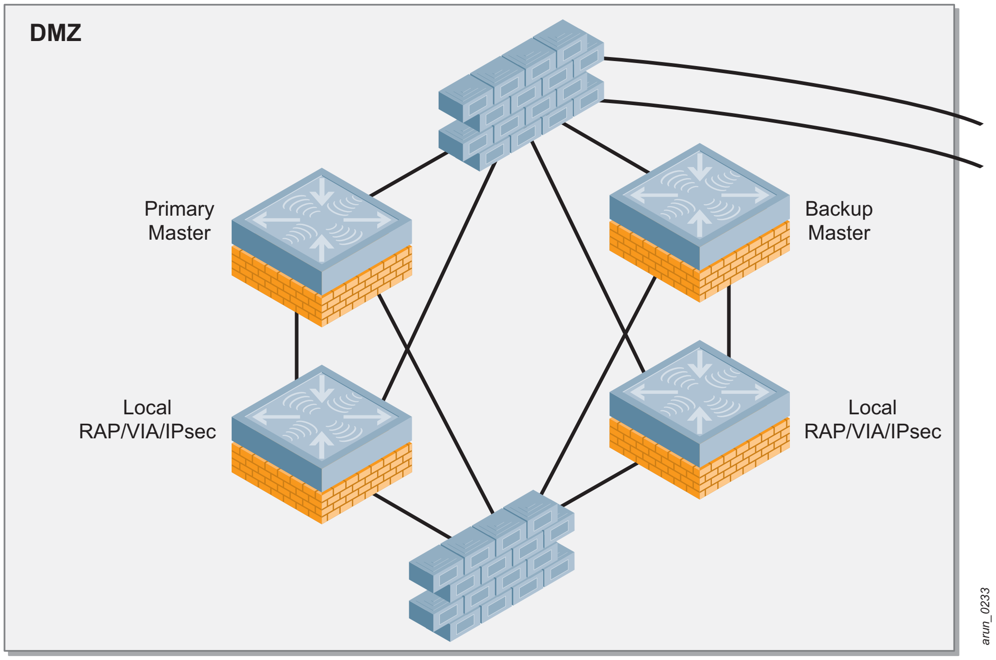

AirWave Deployments

The AirWave Wireless Management Suite (AWMS) is the top tier of management in an Aruba WLAN system, and it runs on a server appliance. This device is always outside the data path, so Aruba recommends that the AWMS appliance be deployed in the data center.

Local Data Center

When deployed locally, no special actions must be taken other than to ensure that the AWMS can reach the mobility controllers it is managing. See the *AirWave Wireless Management Suite 7.0 User Guide* for a complete list of ports that must be open between the AWMS server and Aruba WLAN components.

Remote Data Center

If the AWMS system is going to monitor mobility controllers at remote sites, a VPN connection must be established between the sites to allow AWMS to monitor the mobility controllers securely.

Figure 40 *AWMS access to a remote branch office controller via a site-to-site VPN connection*

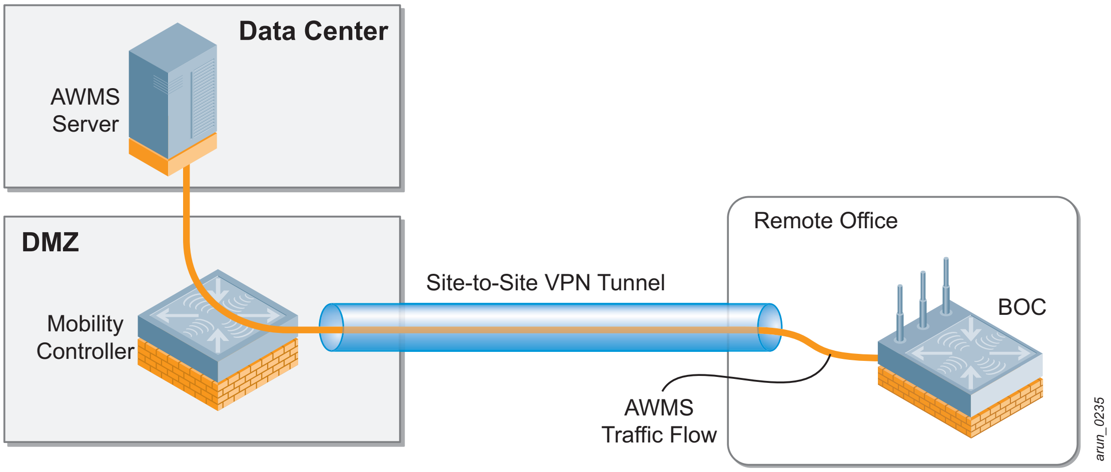

Physical Deployment Considerations and Recommendations

Considerations for Local Mobility Controller Physical Deployments

The following table lists the pros and cons of the different deployment locations. This table covers the local mobility controller, because the master and AWMS server should be deployed only in the data center.

Table 16 *Physical Deployment Considerations*

Location	Pro	Con
Data center	• Centralized location with protected access • Encryption edge-to-core in tunnel mode • Simplified network if Internet access is also through the data center	• May require a “U-turn” if traffic is not destined for the data center

Table 16 *Physical Deployment Considerations (Continued)*

Location	Pro	Con
Core	● Fast connections ● Works well when multiple data centers and Internet access are not co-located	● Potentially limited port interfaces for devices ● Lack of end-to-end encryption for high security organizations ● Often organizations resist placing "bump-in-the-line" devices in the core
Distribution	● Good balance between fast access to data center(s) and Internet ● Ideal for organizations with multiple data centers on a single campus ● Leverage existing redundancy in the distribution layer ● Common deployment location for services infrastructure	● Lack of end-to-end encryption for high security organizations
Access	● Offload at the edge from the mobility controller ● No need for additional PoE switches at the edge	● High cost to provide enough mobility controllers to power APs in large numbers vs. switched infrastructure with PoE ● No redundancy possible without additional switching infrastructure
DMZ	● Ideal placement for locals serving as secure gateways for RAP, VIA clients, IPsec clients, and site-to-site VPNs ● Master mobility controllers may also be placed here to separate remote deployments from the campus ● Can be used for GRE termination from internal switches to force guests out through the DMZ	● Higher security often means separate change control procedures must be followed

Physical Deployment Recommendation

The following table summarizes the Aruba recommendations for where the mobility controller, AP, and AWMS should be placed.

Table 17 *Physical Deployment Recommendations*

Device	Deployment Recommendation
Master mobility controller	Master mobility controllers should be placed in the organization's data center. For organizations that have very small or very large remote access needs, masters may also be deployed in the DMZ.
Local mobility controllers	Local mobility controllers should be placed as close to the destination of the majority of user traffic as possible. This location is commonly either in the distribution layer (more Internet-bound traffic) or in the data center itself.
APs, AMs, wired APs, RAPs	These devices should be deployed at the access layer, where they are accessible for user connectivity.
AirWave Wireless Management Suite	The AWMS server appliance should be placed in the network data center.

Chapter 7

Regulatory Compliance for International Deployments

AP radios must meet government regulatory compliance requirements in the countries where they are installed. This regulatory requirement is based on the mobility controller in the Aruba solution, so care must be taken to select the correct controller for the deployment. The geographical areas of the world for regulatory compliance purposes are broken down as follows:

- USA
- Israel
- Rest of World (ROW)

  In addition, certain Aruba Mobility Controller models are prohibited from being shipped to or operated in other countries.

Figure 41 *US and ROW mobility controller management domains*

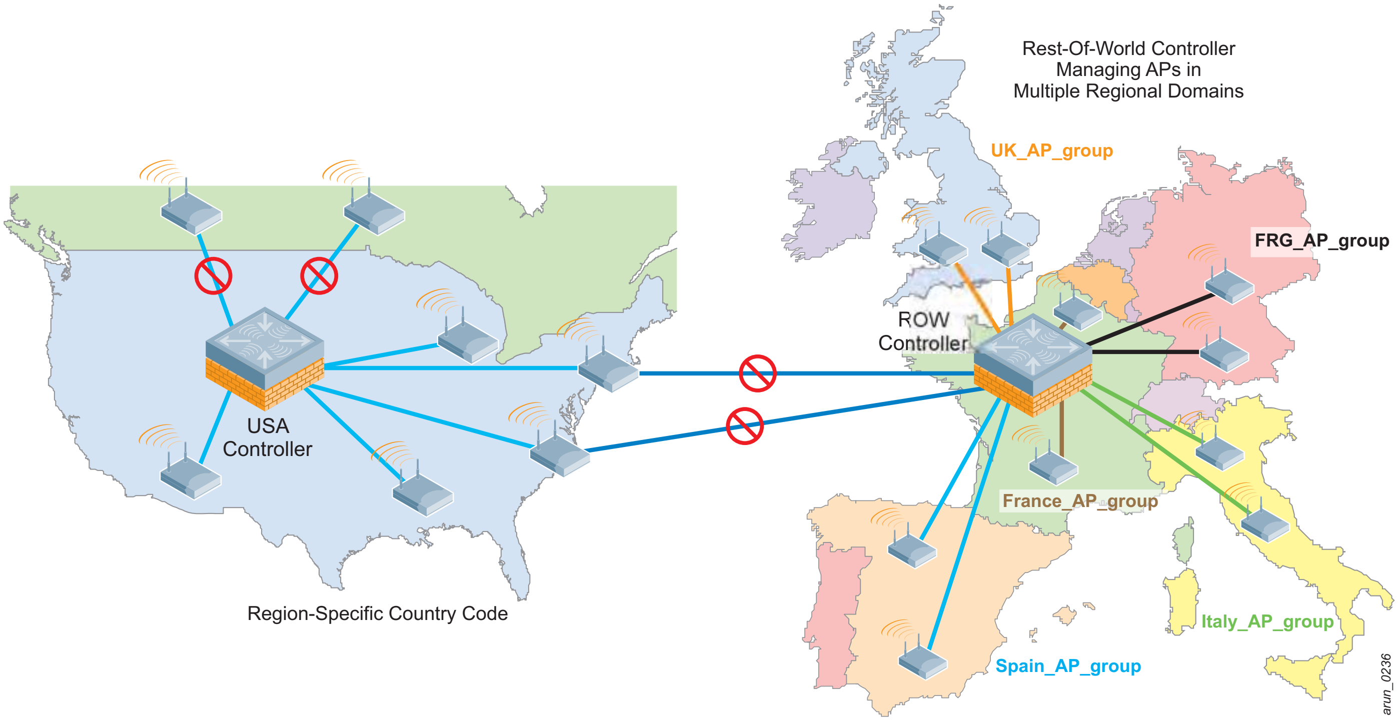

Controller Compliance

When ordering an Aruba Mobility Controller, customers specify a geographic region: United States, Israel, or ROW.

Aruba Mobility Controllers sold in the United States or Israel are physically restricted from managing mobility controllers, APs, and RAPs in other regulatory domains. Administrators cannot assign another regulatory domain to the mobility controllers, APs, or RAPs that terminate at these controllers. However, a ROW mobility controller can properly manage mobility controllers, APs, and RAPs from any unrestricted country and enforce the correct regulatory radio rules.

For example, a US-based mobility controller may not terminate or manage mobility controllers, APs, or RAPs based in Canada or Mexico, nor can it fail over using VRRP to a non-US mobility controller. But a

ROW mobility controller may fail over to an identically configured ROW mobility controller for redundancy purposes.

Figure 42 *ROW mobility controller managing APs in multiple domains*

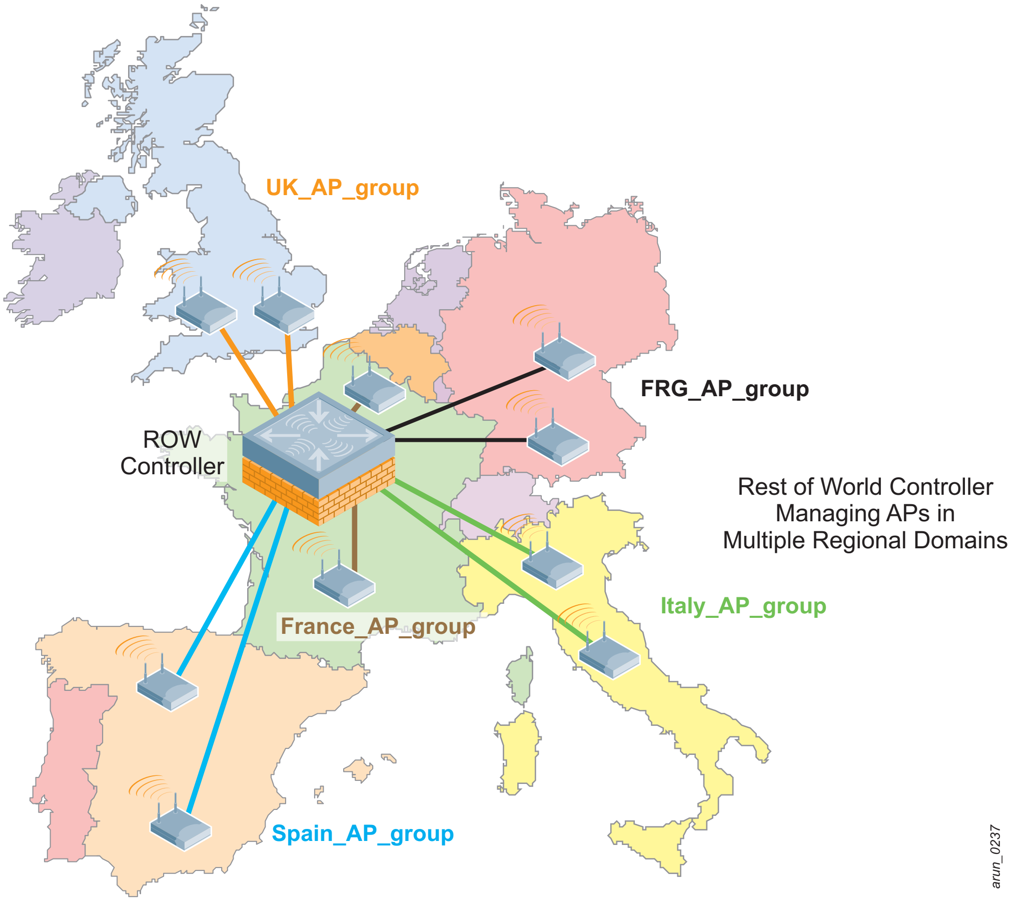

A single ROW Aruba Mobility Controller can manage mobility controllers, APs, and RAPs in France, Germany, Italy, and Spain as long as the APs in each country are properly assigned to separate AP groups. Each AP group must be assigned an RF management profile with the correct country code that corresponds to the physical location of the APs.

Recommendations for International Deployments

Use this checklist to verify that your Aruba design complies with the host country laws and regulations:

1. Review all mobility controllers that participate in VRRP clusters to confirm that all models have identical country SKUs.
2. Review all mobility controllers, APs, and RAPs that terminate on US-based mobility controllers and make sure that they are all in the US.
3. Review all mobility controllers, APs, and RAPs that terminate on Israel-based mobility controllers and verify that they are all in Israel.
4. Make lists of all mobility controllers, APs, or RAPs by country to create the proper Regulatory Domain profiles when you configure the system.
5. Purchase any additional mobility controllers necessary to achieve regulatory compliance.

Chapter 8

Logical Deployment Considerations

The Aruba Mobility Controllers, APs, and AirWave management platform fall into four logical layers: management, network services, aggregation, and network access.

Figure 43 *Logical deployment overview*

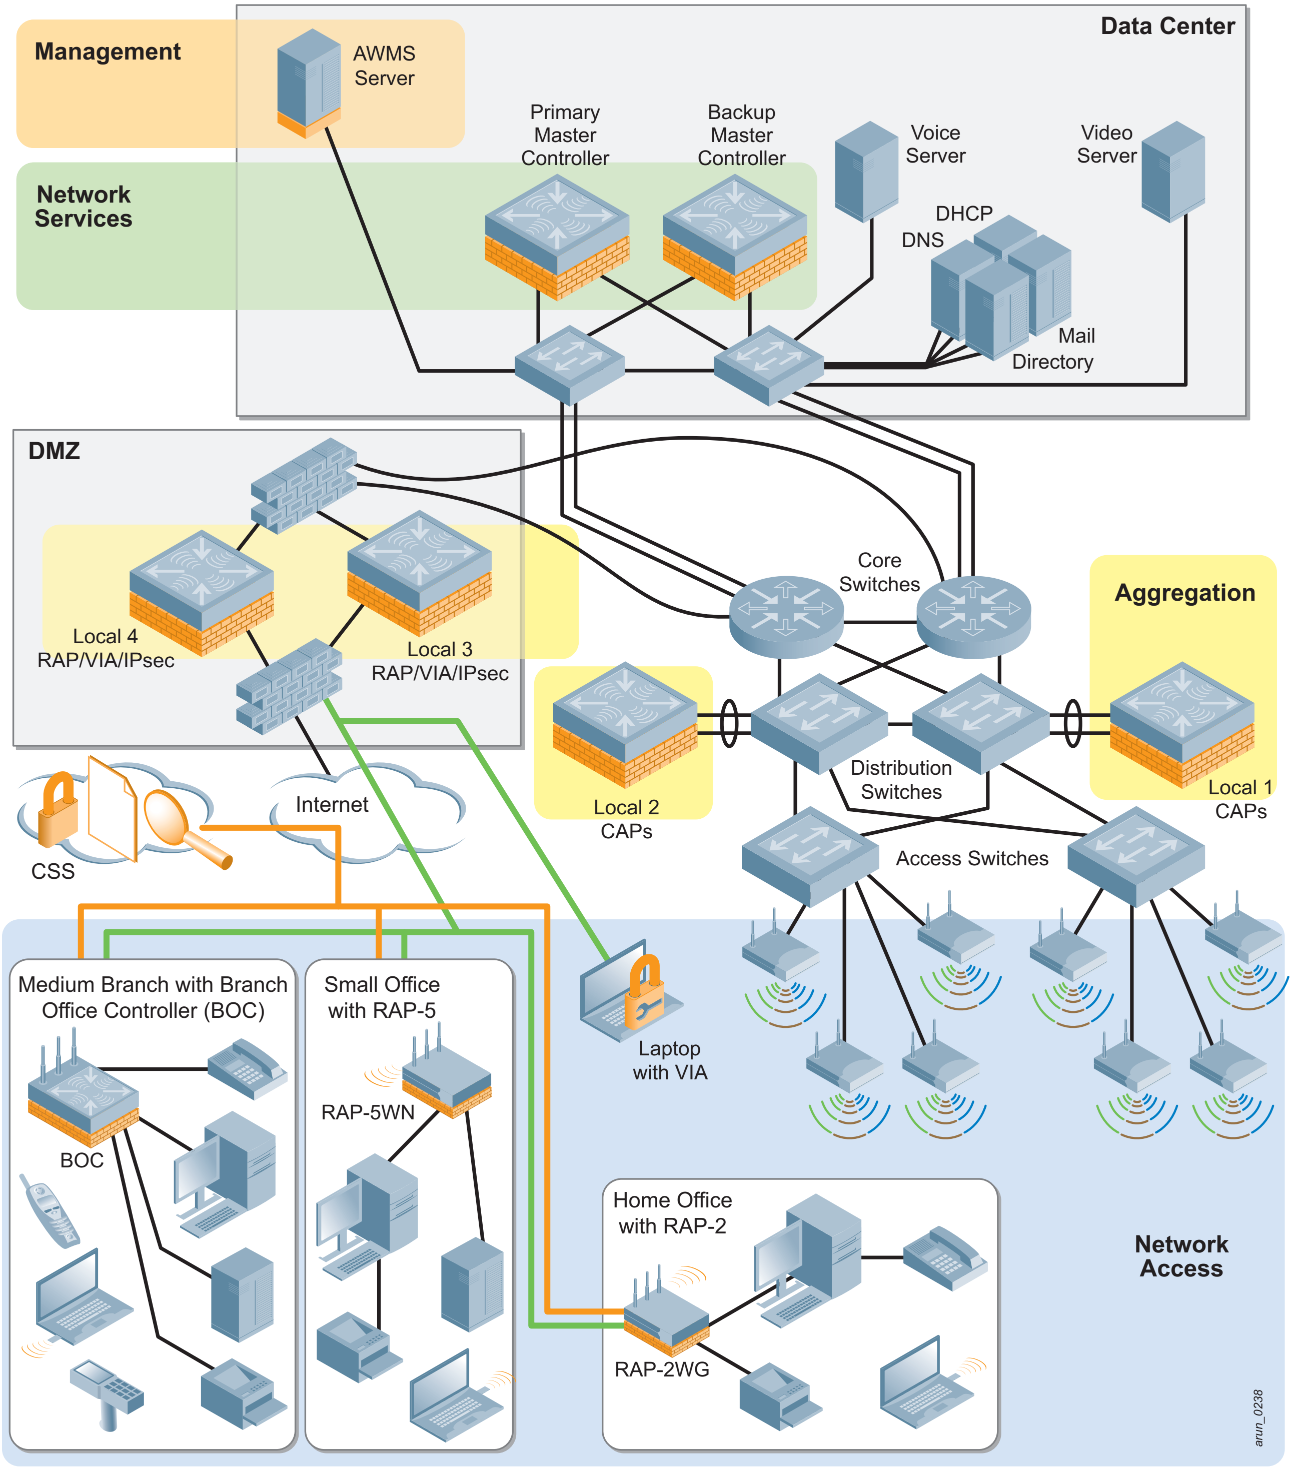

- **Management:** The management layer consists of the AWMS server. The AWMS provides a single point of management for the WLAN, including reporting, heat maps, centralized configuration, and troubleshooting.
- **Network Services:** The network services layer provides a control plane for the Aruba system that spans the physical geography of the wired network. This layer consists of master mobility controllers. The control plane does not directly deal with user traffic or APs, but instead it provides services such as white list coordination, valid AP lists, CPsec certificates, WIP coordination, and RADIUS or AAA proxy.
- **Aggregation:** The aggregation layer is the interconnect point where the AP, AM, wired AP, and RAP traffic that is destined for the enterprise network is aggregated. This layer provides a logical point for enforcement of roles and policies on centralized traffic that enters or exits the enterprise LAN.
- **Network Access:** The network access layer is comprised of APs, AMs, wired APs, RAPs, and physical controller ports that work together with the aggregation layer controllers to overlay the Aruba system. When policy-based or bridge forwarding modes are used, firewall policies are applied at AP. Bridge mode traffic never reaches the controller, while split-tunnel traffic is forwarded only to the aggregation layer for enterprise destinations and traffic not directly bridged.

User VLANs

The VLANs that support user traffic and that the client device uses to receive its IP addressing information are not always that same as the VLAN that the AP is plugged in to. In many cases, the user VLAN has nothing to do with the VLAN that the AP is connecting through. The user VLAN assignment varies depending on forwarding mode, so each forwarding mode is examined.

User VLANs in Tunnel and Decrypt-Tunnel Modes

In the tunnel and decrypt-tunnel forwarding modes, user traffic flows transparently across the network in a GRE tunnel. The user traffic is not converted to an Ethernet frame and placed in a VLAN until it reaches the mobility controller. The user VLAN does not exist at the AP that is providing access, so the VLAN the user is actually placed into does not need to exist there either.

The wireless traffic is processed at the AP, but because it is sent over a GRE tunnel, the user does not need to be in the same VLAN as the AP. The easiest way to think of this construct is that the user essentially is connected directly to the mobility controller. IP addressing is based off of a logical design for the user, as opposed to the physical port that the AP is plugged in to, as is the case with bridge mode APs. The following figure shows an AP attached to an edge switch with a VLAN that extends to the mobility controller.

Figure 44 *AP plugged into a local switch, accessing the mobility controller*

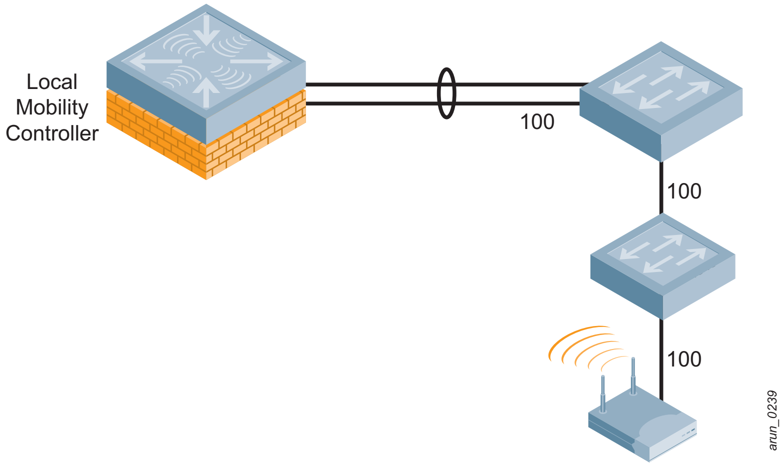

In this case, the VLANs that the users are assigned to do not exist at the AP. Those VLANs exist only on the mobility controller itself. This configuration simplifies the edge of the network, because all user VLANs are not required to reside at the edge switches and they need only be trunked to the mobility controller. The following figure shows the actual VLAN of the users, which exists only from the mobility controller through the switch to the router.

Figure 45 *User VLAN, logical connection*

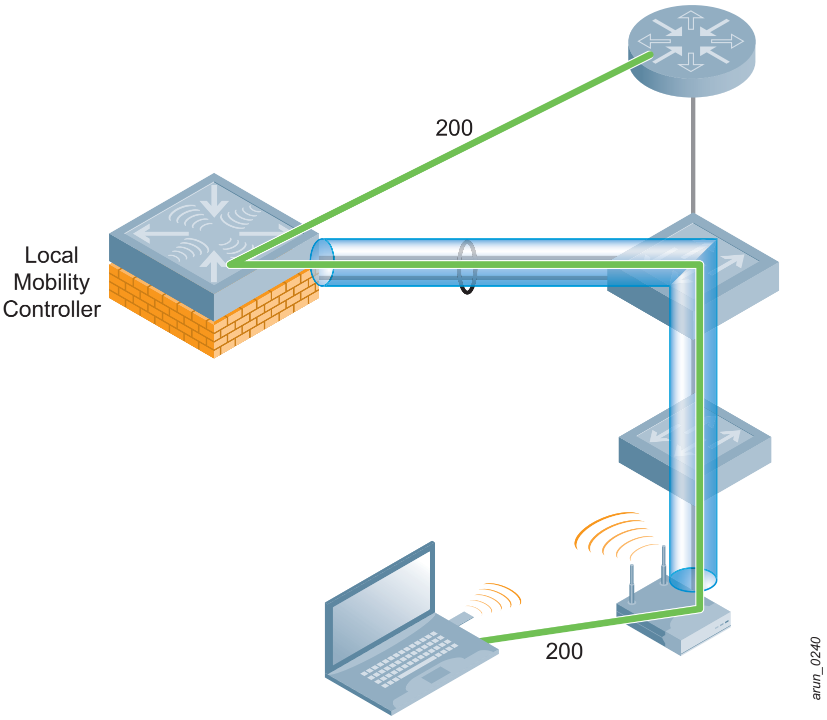

The advantage of this design is a simplification of the network and flexibility of terminating users. When the organization needs additional user VLANs, these can be created only at the switch that connects to the mobility controllers, and no change needs to be made to the APs or the network edge.

User VLANs in CAP Bridge Mode

When APs are used in bridge mode, the user VLAN and the AP VLAN are typically the same VLAN, because this model operates only on a flat Layer 2 network. In this case, the AP is handling the traffic exclusively for the user and bridging it locally instead of sending the traffic back to the mobility controller for processing.

Figure 46 *Users and APs in a bridge mode deployment share the same VLAN*

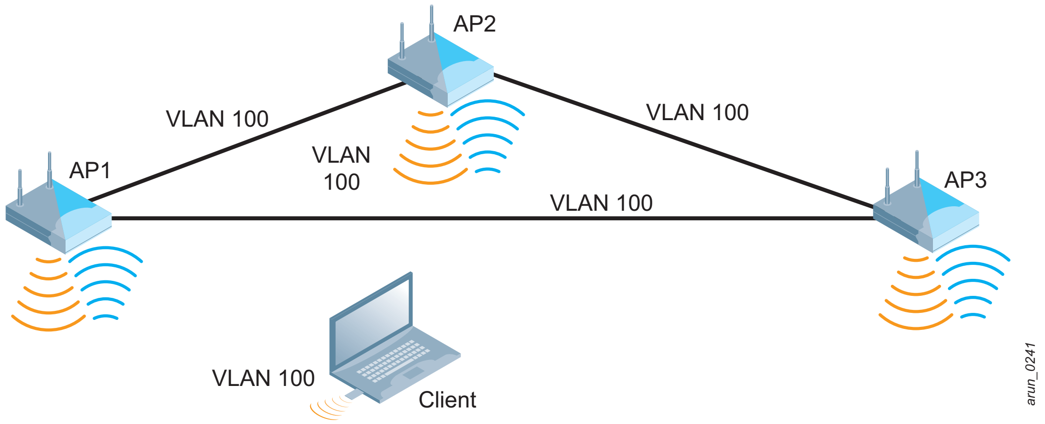

User VLANs in Remote AP Bridge Mode

In RAP mode, the RAP can be configured to act as the local DHCP server for any clients that are attached to bridge mode SSIDs or ports. In this case, a VLAN must be defined on the RAP and the controller, and an associated DHCP pool must be configured. This configuration is pushed down to the RAP and is used by any clients that associate to the RAP on a bridge mode connection. The DHCP scope is local to the RAP itself, so the RAP must perform NAT translation on all traffic leaving the upstream interface just as a typical home router would.

Figure 47 *User VLANs in RAP bridge mode*

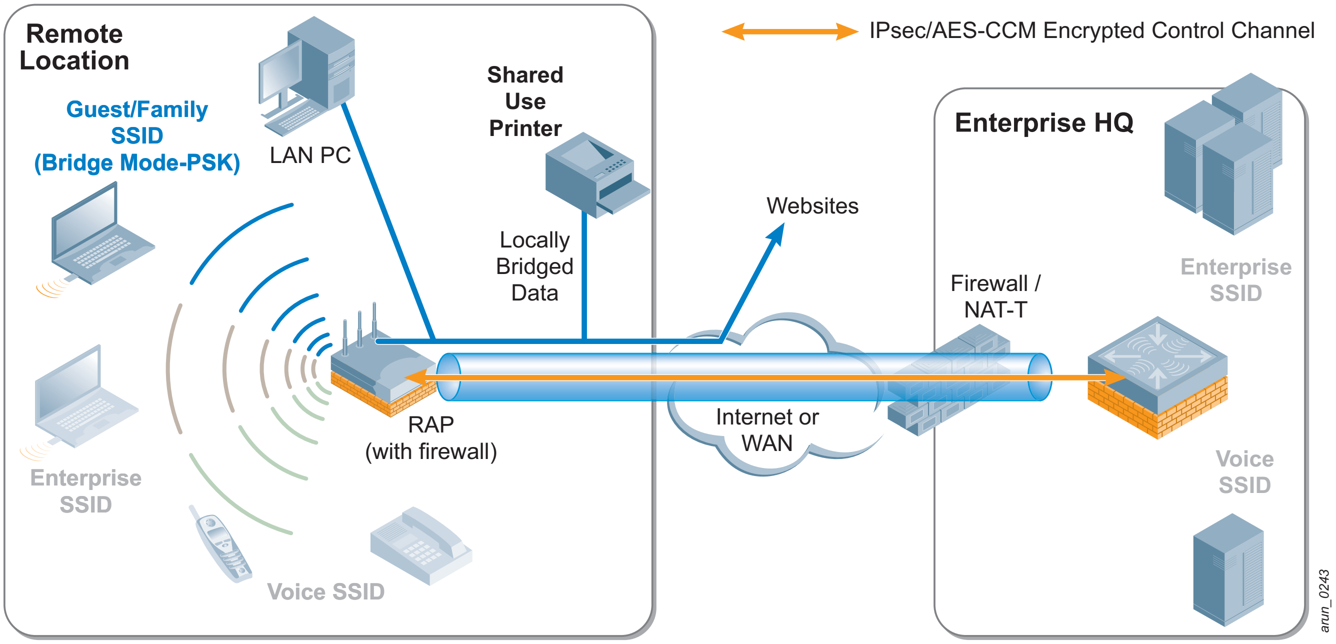

User VLANs in Split-Tunnel Mode

Split-tunnel mode is similar in operation to the tunnel and decrypt-tunnel modes, except that the AP applies firewall policy at the edge and makes routing decisions for the client. IP addressing is supplied from the mobility controller centrally. The AP also exists in a local subnet, though this may not be defined as a VLAN.

Figure 48 *User VLANs in split-tunnel mode*

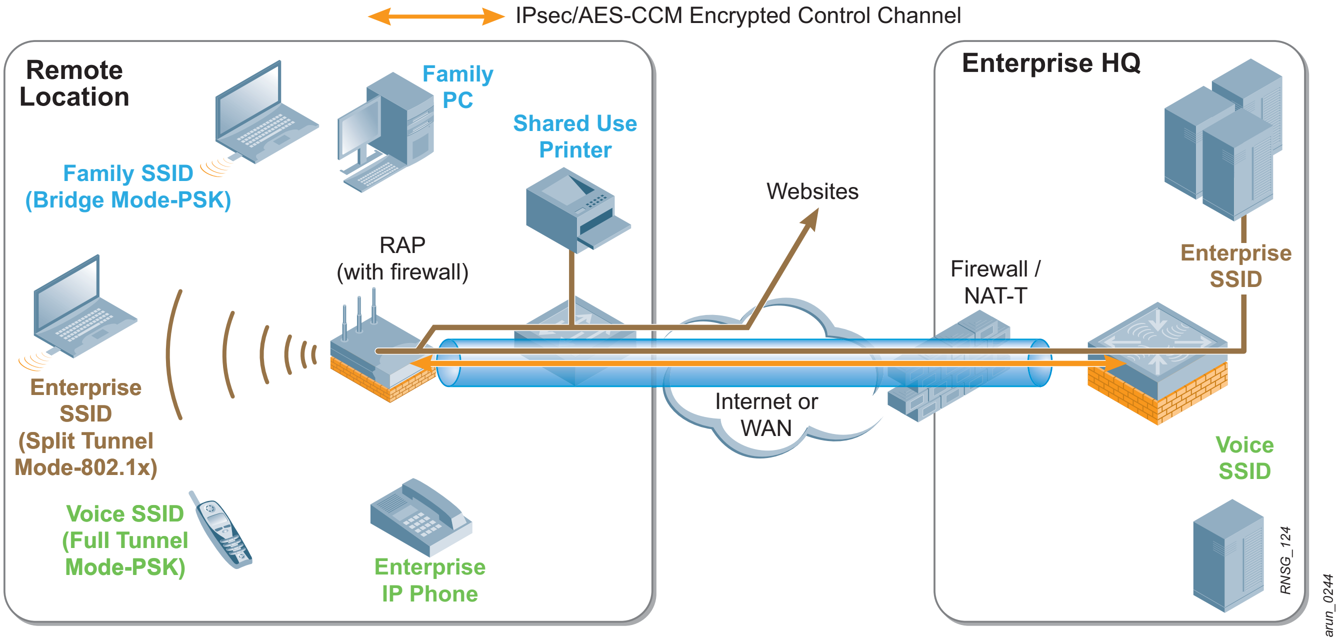

Guest VLANs

Dedicated guest VLANs are common in networks to limit guest access to other parts of the network and they take two forms:

- **VLANs to the DMZ:** Limit guest access by using a management construct.
- **VLANs just at the mobility controller:** Limit guest access by using firewall policy.

When the VLAN is run from the controller to the DMZ, users are placed in this VLAN and sent to the DMZ. Routers in the network forward traffic only to the DMZ and do not allow the users to route to other VLANs. This action protects the local infrastructure if the VLAN design is secure, but it does nothing to stop users from interacting with one another on the same guest VLAN.

Figure 49 *Guest VLANs without firewall enforcement*

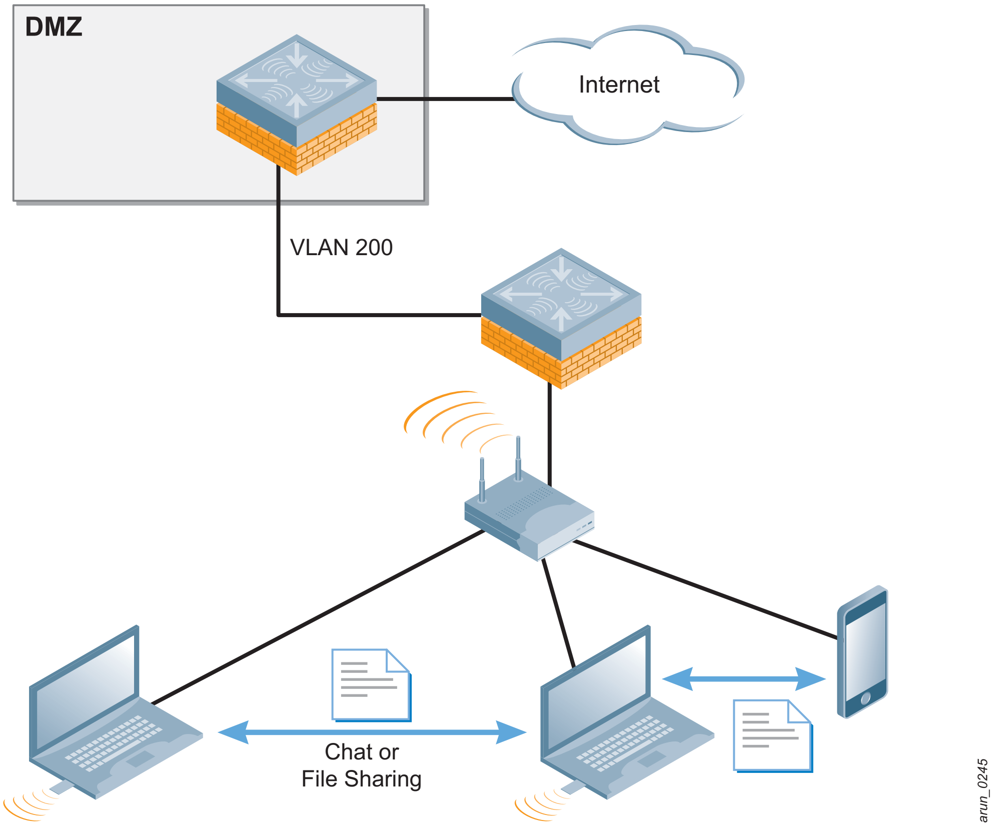

When the Aruba PEF is used, the guest VLAN typically exists only on the local mobility controller. The local will act as the DHCP server, and the firewall policy is used to limit user traffic. A typical policy will allow the user to receive DHCP and DNS from the local network, but will then prevent all other traffic destined to the local network and allow only Internet access. Typically guest users in this scenario receive a private, nonroutable IP address and NAT is performed as their traffic leaves the controller on a public VLAN.

Figure 50 *Guest VLANs with firewall blocking traffic between devices, allowing Internet access*

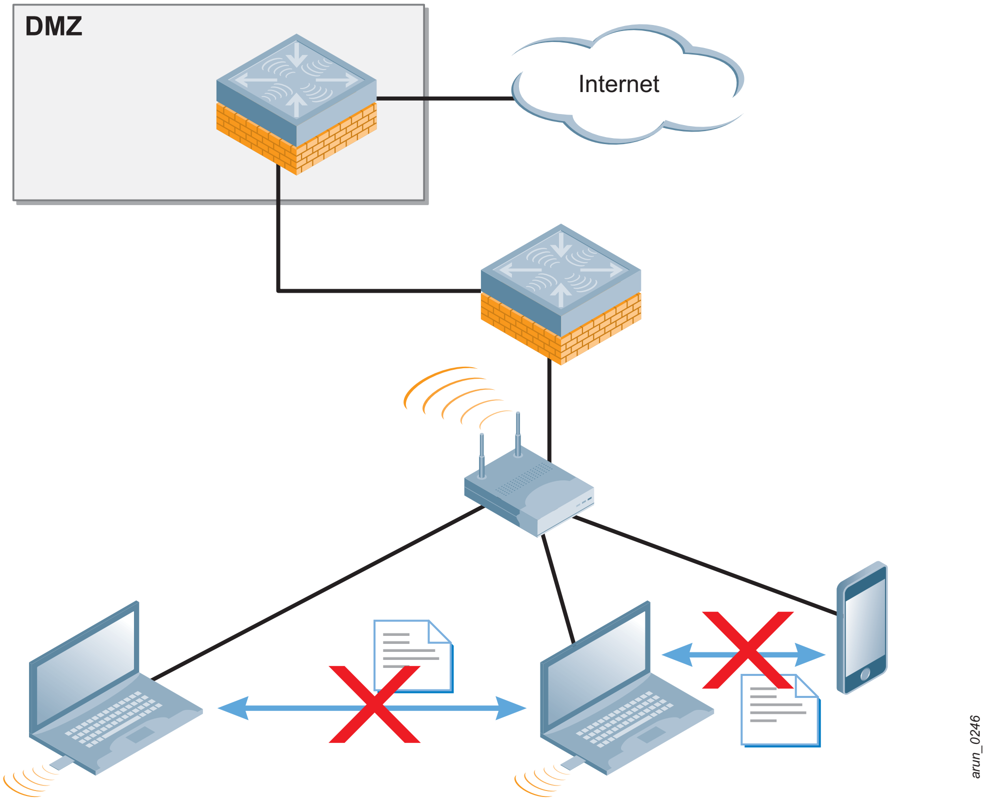

These two delivery mechanisms are not a "one or the other" decision, and they can be combined. Aruba recommends that role-based firewall policies be applied to guest users even when using a dedicated VLAN that is routed to the DMZ. For more security, users may want to use GRE instead of a VLAN to force clients to the DMZ controller.

Figure 51 *Guest VLANs using firewall to prevent device communication, GRE to external DMZ controller*

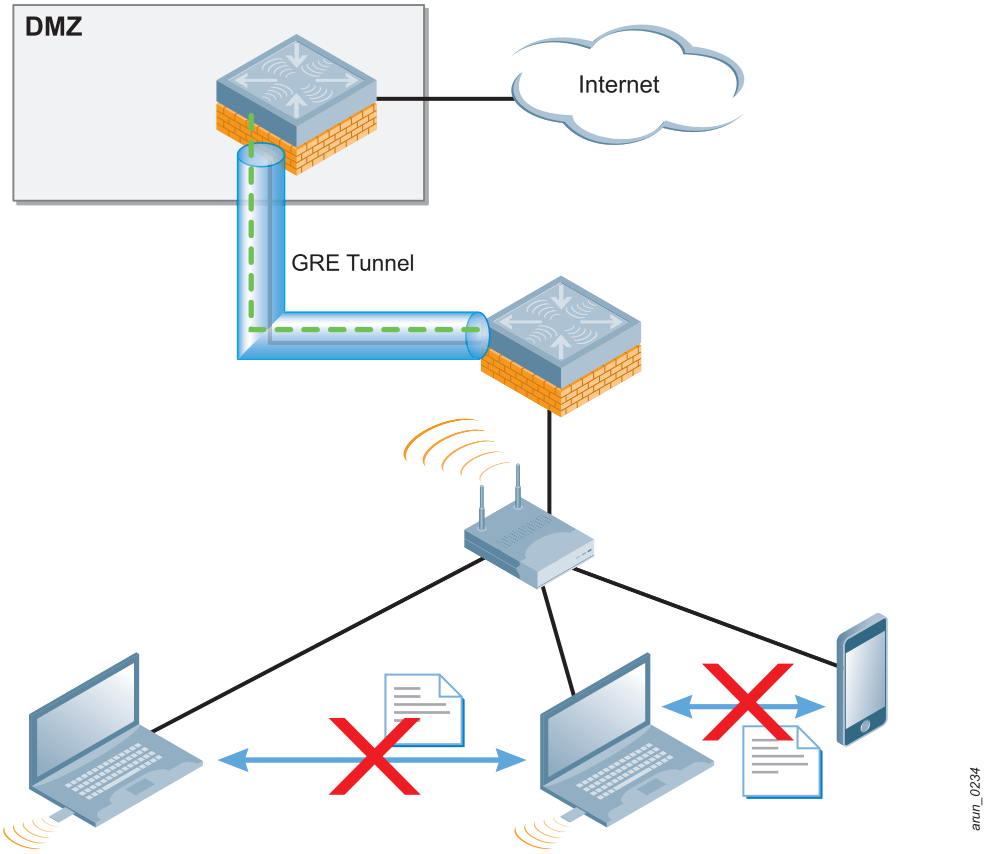

Dedicated AP VLANs

Early in the deployment of wireless networks, dedicated AP VLANs were used to segregate the wireless traffic from other wired traffic. This segregation was done to force wireless traffic through firewalls and IPsec concentrators to secure wireless connections after WEP was broken.

Figure 52 *Historical AP VLAN model*

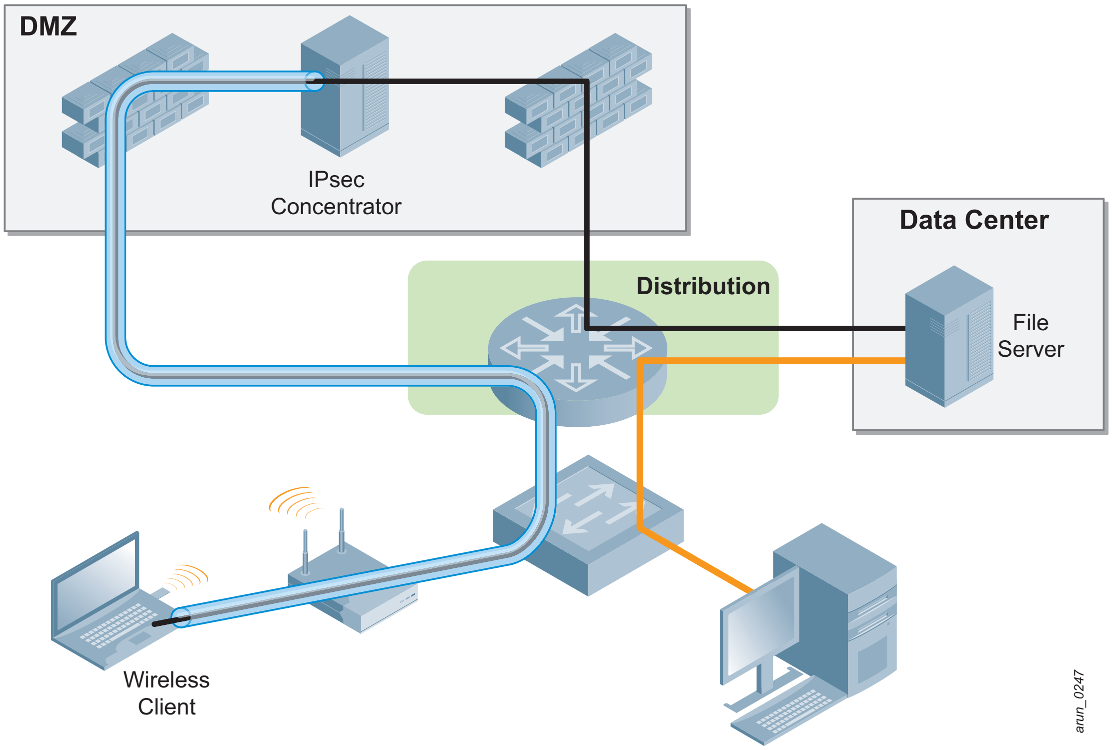

This method not only leads to management overhead to ensure that each AP is plugged into an "AP port" on the switch, but also that the switch is configured correctly. The other downside to this approach is that AMs become less effective, because they can no longer see user traffic that may be exiting a rogue AP on the wired side of the network.

Figure 53 *802.1X, user passes, AP fails and is assigned a VLAN t hat is only routable to the mobility controller*

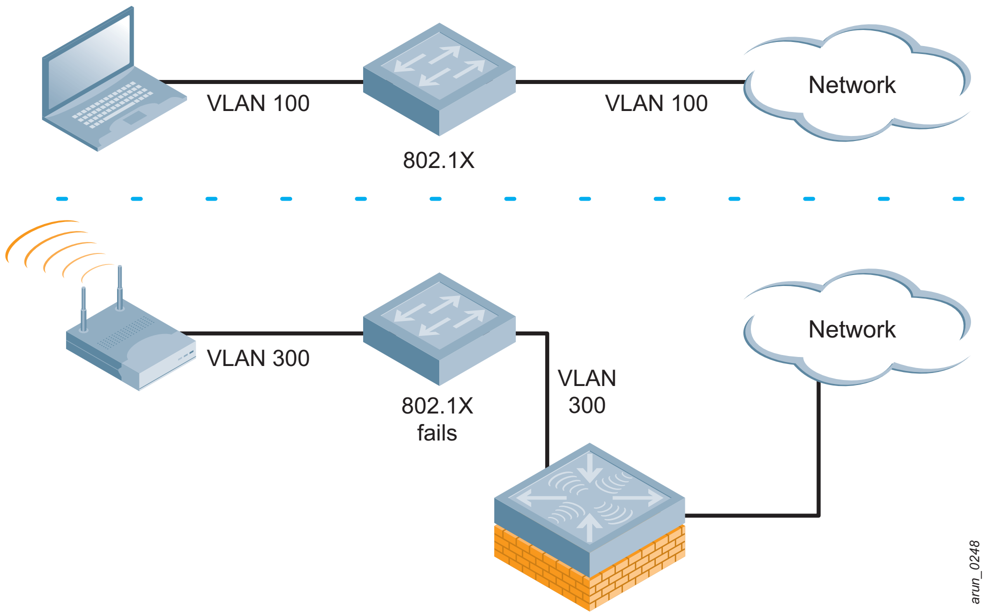

The only recommended use for a separate AP VLAN is in networks where 802.1X is configured on the edge switch to do link layer authentication of users. The AP does not support an 802.1X supplicant, so it is recommended that the wired switch be configured to place "failed" devices in a special AP VLAN that only is only routable to the mobility controller.

Quarantine VLANs

Quarantine VLANs are common in networks where network access control (NAC) has been integrated. This VLAN is used for devices that have failed their health check and are put into the quarantine VLAN until the device can be brought in-line with policy. The problem with this traditional method is that a set of infected stations in the same VLAN tends to lead to more infections. If the users are able to remediate, they must then be moved back to the "production" VLAN and receive a new IP address.

Figure 54 *Quarantine VLAN does nothing to stop cross-device infection*

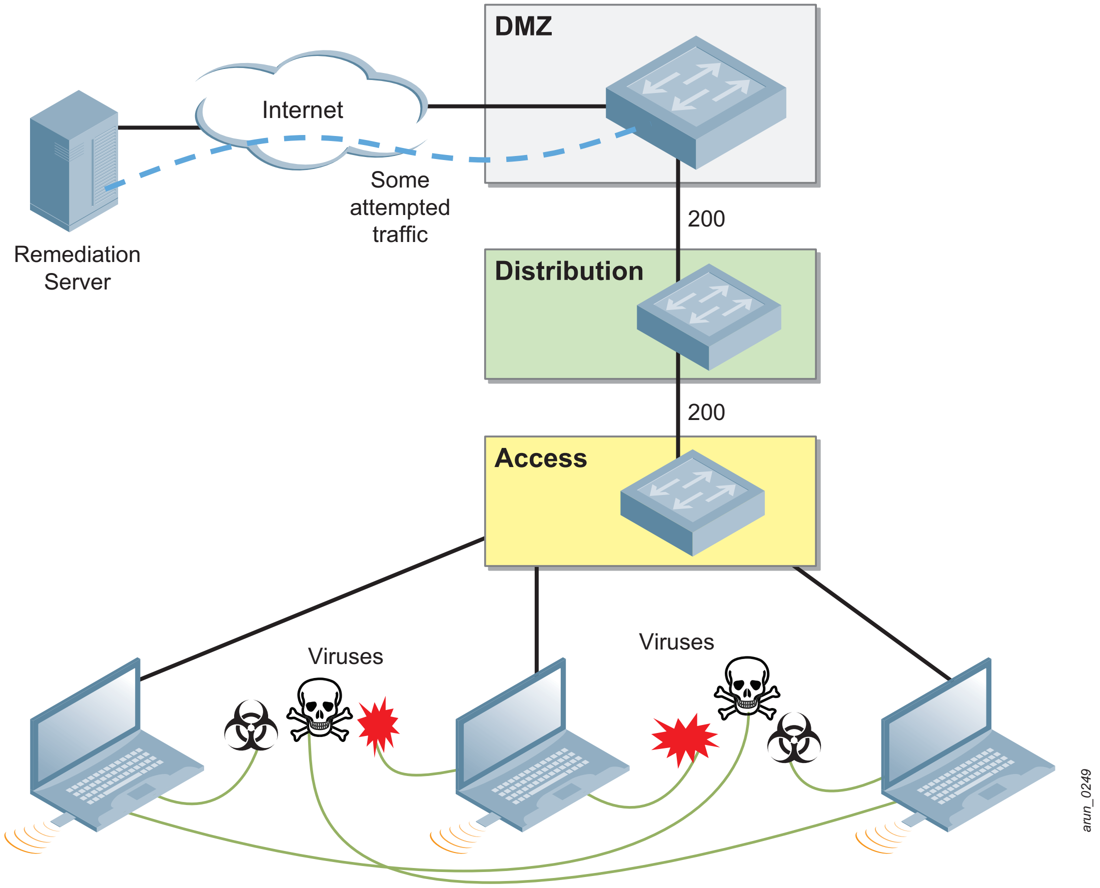

Instead, Aruba recommends that the PEF-NG firewall be used to put users into a quarantine role. This role should allow stations to access only remediation resources, either locally or on the Internet. These resources could include anti-virus vendors, operating system vendors, and software vendors. All other traffic should be denied, which removes the ability of the station to infect other users. After the station is remediated, it does not need to reboot or renew its IP address because the station does not switch VLANs. The user simply is placed in the production role and allowed to use the network fully.

Figure 55 *Using the firewall to limit the spread of viruses in the remediation role*

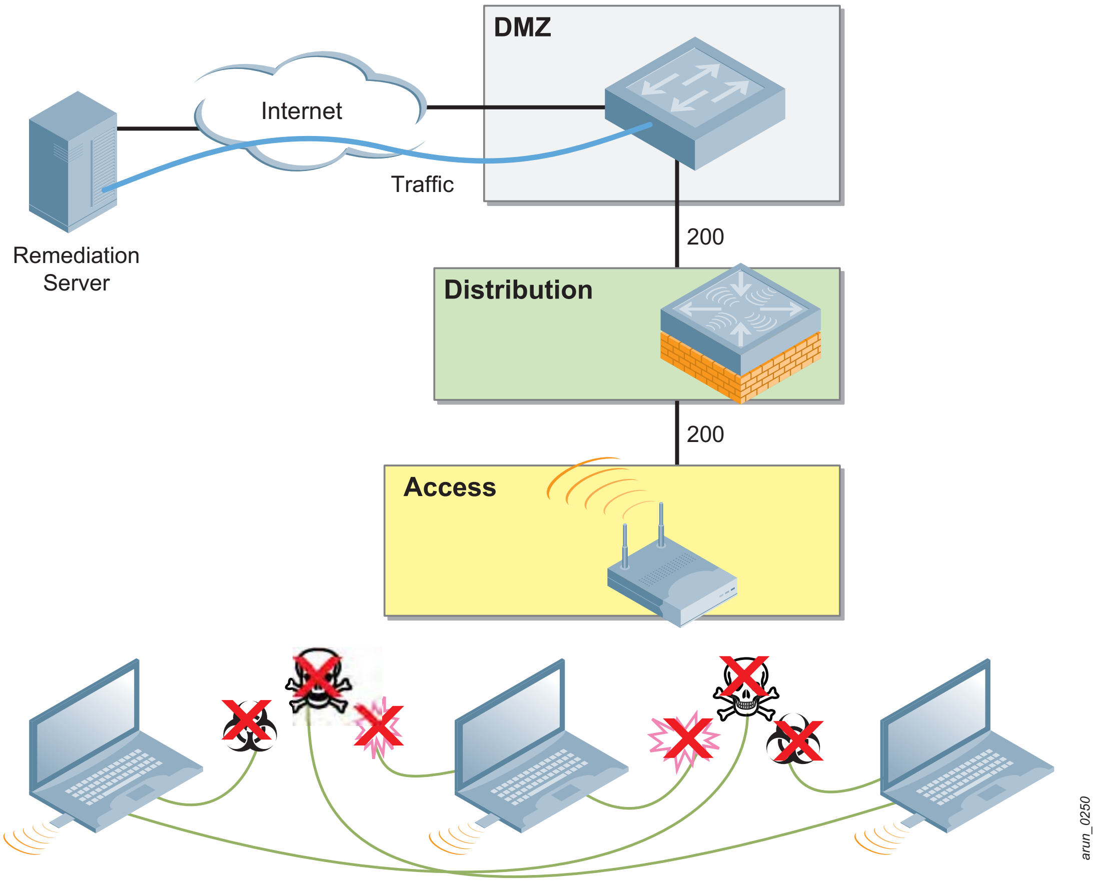

VLAN Pools

Network administrators prefer to keep subnet sizes down to a Class C size network. This network has a subnet mask of /24, which yields up to 253 user devices per subnet. This size is considered manageable and helps to limit the broadcast domain size. In networks where this subdivision needs to be logical as opposed to physical, VLANs are employed to limit broadcast domain size. The issue arises when enough users exist to exceed a single subnet, which is a common occurrence because the WLAN has gone from a convenience network to a part of the critical network infrastructure.

The traditional methodology for dividing up large groups of wireless users is to place a set of APs in a VLAN and have all users associated with those APs placed into that single VLAN. This method works if the user count never goes above the subnet user count limit and if users have no need to roam outside of the AP group. This method limits the size of a subnet, and it is typically deployed only in small networks with a single subnet.

Figure 56 *VLANs spread across groups of APs*

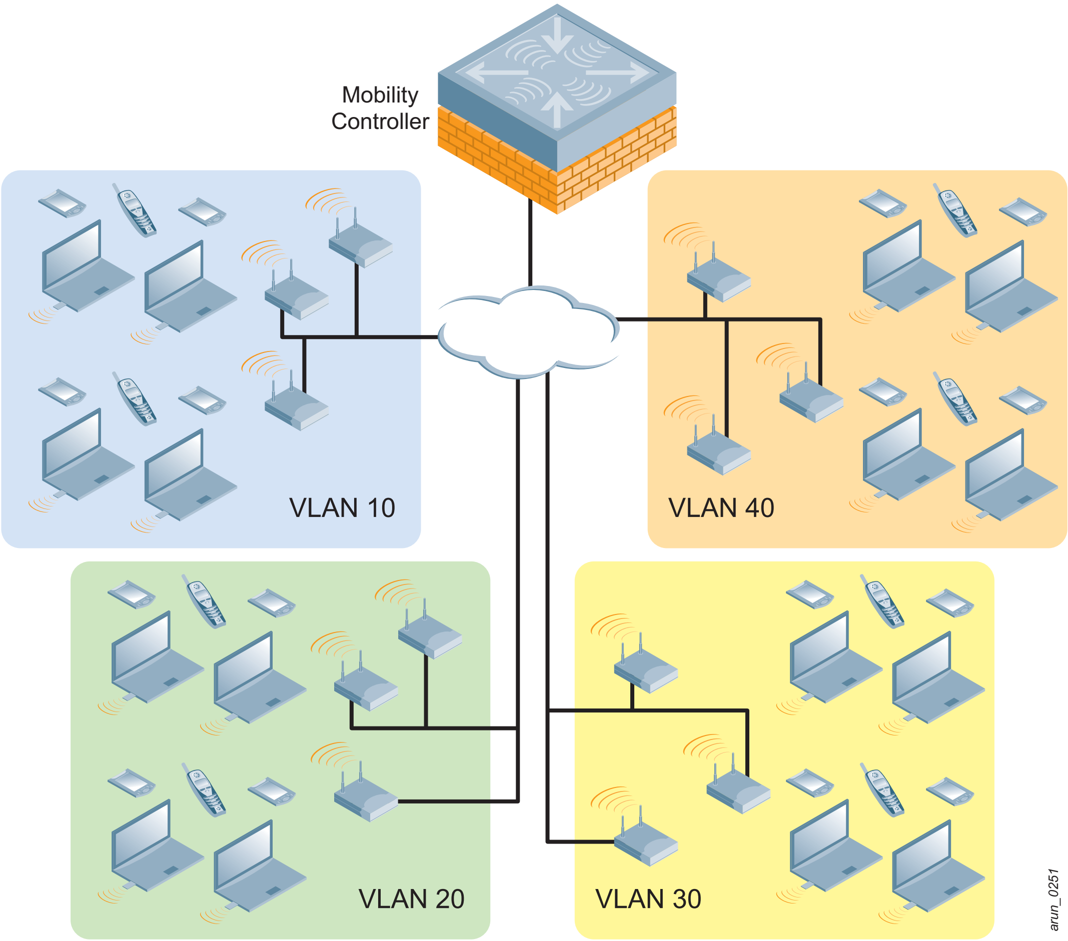

However, this method tends to fail when large groups of users need to meet in a single location like a lecture hall, or an “all hands” meeting, or where roaming across APs is likely to occur.

The Aruba VLAN Pooling feature allows a set of VLANs to be assigned to a designated set of virtual APs. These VLANs can be configured as a noncontiguous set, a contiguous range, or a combination of the two. As an example, the set could be VLAN numbers 10, 20, and 30. The set could also be VLAN numbers 2 through 5. These methods can be combined to provide a set such as 3, 5, and 7 through 10. This flexibility allows you to assign users to VLANs that may already exist in the enterprise. VLAN pools are the method that Aruba recommends for handling user VLANs any time two or more user VLANs are needed to handle the user load from a single set of APs going to a single mobility controller.

Figure 57 *VLAN pools distribute users across VLANs*

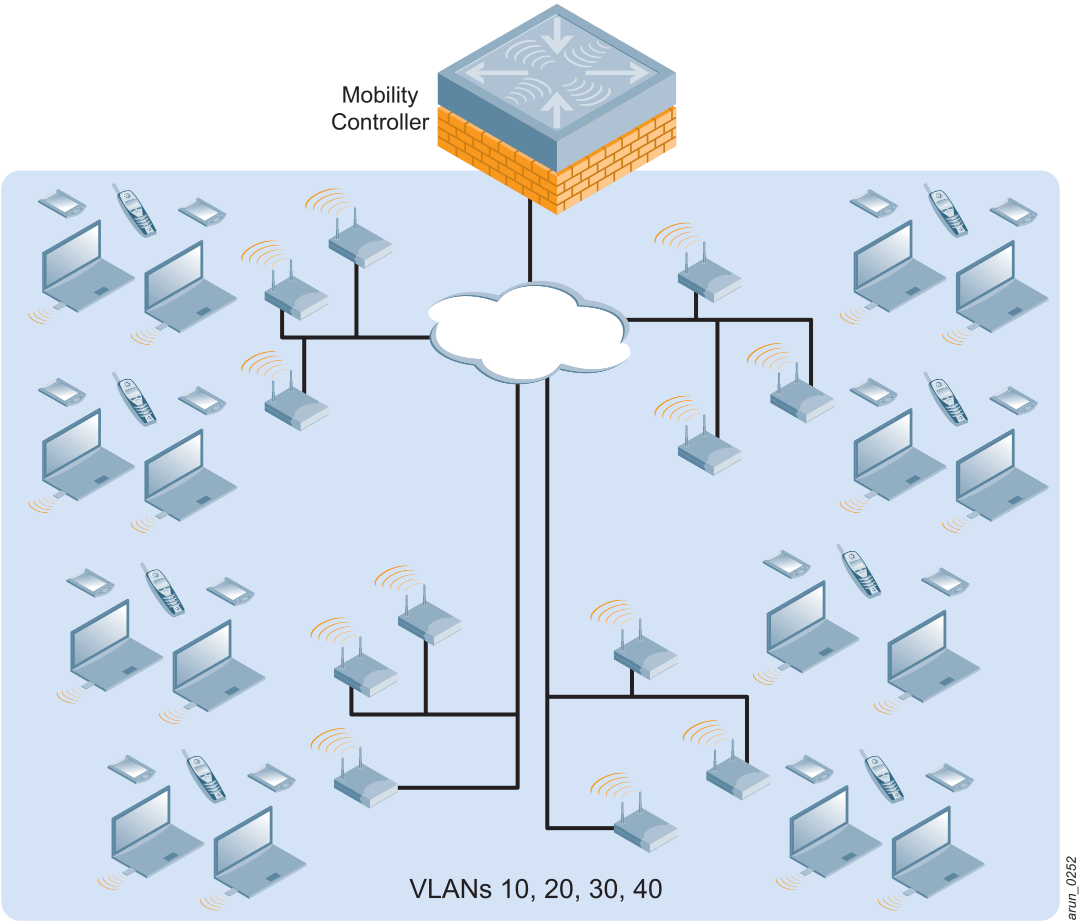

The system works by placing users in one of the VLANs in the pool. VLAN placement is determined using the user MAC address and running it through a hash algorithm. The output of this algorithm places the user into one of the VLANs in the pool and ensures that the user is always placed into the same pool during a roaming event. As the user associates with the next AP, their address is hashed and they are placed into the same VLAN on the new AP. The user can continue to use their existing IP address with no break in their user sessions. This feature requires that the same VLAN pools be deployed on all controllers that will service these clients.

Packet Sizing

Wherever possible the network should be configured to support jumbo frames to avoid fragmentation of the packets. When this configuration is not possible due to lack of hardware support, the maximum MTU should be configured on all devices. This setting is especially important for video transmissions, because the loss of key video frames can cause the entire packet to be retransmitted. It is also important that transmitters are aware of the frame size limitations. For instance, video servers should be configured to use the maximum MTU of the network to limit their packet size to the network-supported maximum and avoid fragmentation.

Default Gateways and Routes

Users terminate on the Aruba Mobility Controller, so there can be some debate about where the default gateway should exist and how routing table updates should occur. This section describes the options for deploying the mobility controller as the default Layer 3 gateway as opposed to a Layer 2 device. Also discussed is how user subnets should be routed if the controller is selected as the gateway.

Layer 2 Deployments

In a Layer 2 deployment, the mobility controller is a "bump in the line" for user traffic. Wireless sessions are inspected by the firewall and forwarded to the appropriate VLAN, but the mobility controller is not the default gateway. This deployment model is typically used in campus networks where an existing Layer 3 switch is already functioning as the default gateway and makes routing decisions for the network. This deployment model is recommended where multicast routing will occur.

Figure 58 *Mobility controller in a Layer 2 deployment*

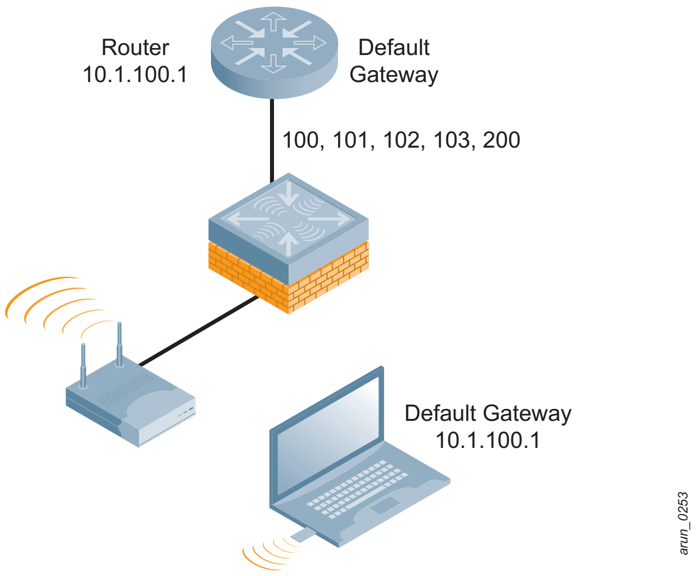

Layer 3 Deployments

The other alternative is a Layer 3 deployment where the mobility controller is the default gateway for the subnet. This deployment is common for remote networking, where the users receive their IP addressing from the mobility controller, and for site-to-site VPN applications to the branch office.

Figure 59 *Mobility controller as the default gateway*

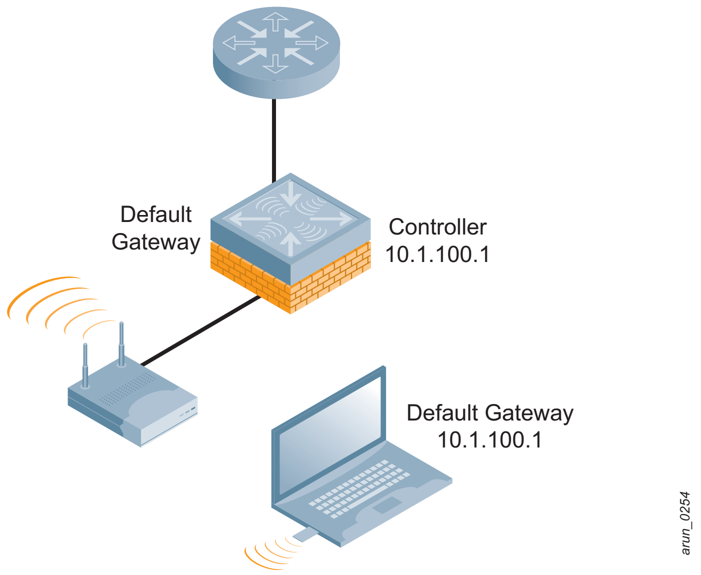

When a RAP is deployed, all addressing is delivered from the mobility controller to the client machines. Machines on the same site may receive different addresses from different pools, which would make routing difficult for a traditional routed network to manage. The mobility controller is the one vending the addressing, so it is logical for it to also act as the default gateway for those subnets.

Figure 60 *Mobility controller as the default gateway for branch offices*

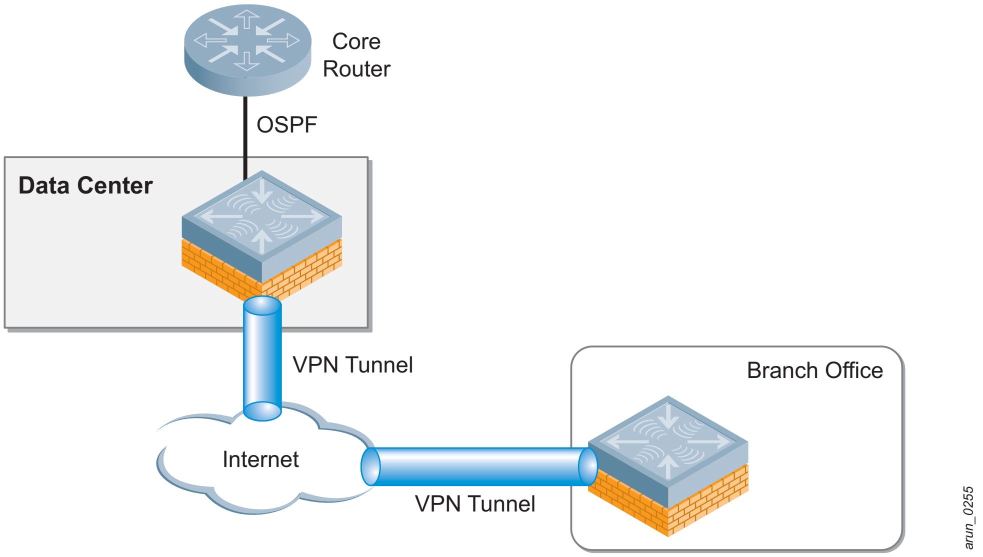

The mobility controller can also terminate VPN sessions, including site-to-site VPNs from other mobility controllers in branch office deployments. In these cases, the local branch typically has a DHCP server local to the site. The mobility controller that establishes the VPN connection is the default router for that site. The mobility controller that acts as the VPN head end will be the gateway for the rest of the network to the branch office.

Figure 61 *Mobility controller providing guest services as the Layer 3 gateway*

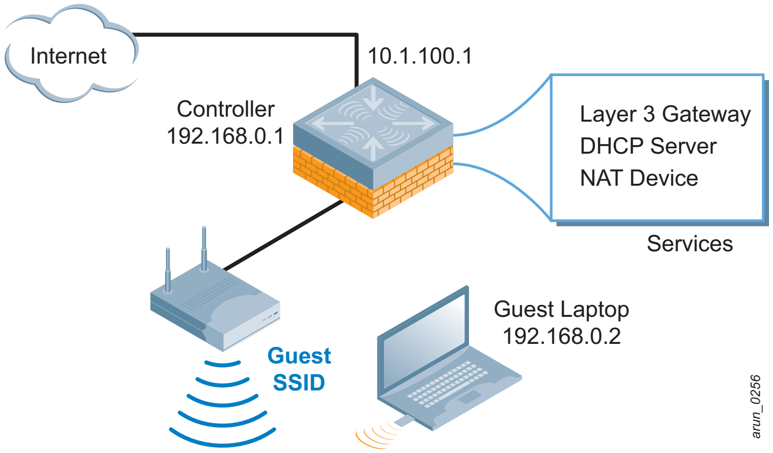

The final case where the controller is typically the Layer 3 device is when it exists as the default router for a nonroutable guest network. When a guest network is deployed in private IP space and is not routable from the general network, the mobility controller is normally configured to act as both the DHCP server and NAT device for the guests.

Static Routes and OSPF

When the controller is deployed as the default gateway for a particular subnet, routers in the network need to know how to reach that gateway. The two methods for handling these advertisements are static routes and dynamic routing protocols. Network managers prefer to avoid static routes where possible, because any change to the network topology requires an update to the static routing table. When dynamic routing protocols are used, no manual updates must be made to routing tables.

Figure 62 *OSPF running between routers and the mobility controller*

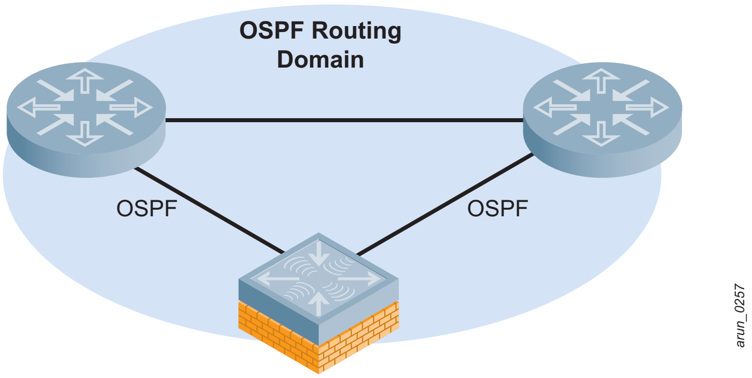

The Aruba Mobility Controller supports running the dynamic routing protocol called Open Shortest Path First (OSPF). The implementation allows the mobility controller to operate in either stub or totally stub mode. This capability allows the mobility controller to advertise its routes into the network without the overhead of maintaining the full routing table.

Logical Design Recommendations

Due to the flexible nature of the Aruba deployment models, logical design recommendations depend on the type of deployment, either campus or remote.

Table 18 *Logical Design Recommendations for Campus and Remote*

Service	Campus	Remote
User VLANs	Use VLAN pools to control subnet size.	Use VLAN pools to control subnet size.
Guest VLANs	Not needed except on the controller, use NAT and PEF-NG to control access.	Not needed except on the controller, use NAT and PEF-NG to control access.
AP VLANs	Do not use dedicated AP VLANs.	Do not use dedicated AP VLANs.
Quarantine VLANs	Not needed, use PEF-NG to control access.	Not needed, use PEF-NG to control access.
Jumbo Frames	Enable jumbo frames if possible, or the largest frame size available. Make sure servers are configured to use the maximum size possible frame to avoid fragmentation.	N/A
Default Gateway	Not for user VLANs. The controller should be the default gateway for guest VLANs.	The controller should be the default gateway for all user subnets.

Campus Logical Design Recommendations

- **User VLANs:** If more than one user VLAN is required, Aruba recommends that VLAN pools be used to distribute users more evenly across the pools. By using multiple VLANs in a VLAN pool, the size of broadcast domains are reduced and the configuration is simplified for the network manager.
- **Guests VLANs:** Though guest VLANs are common in many deployments for historical reasons, guest VLANs that cross the internal network to the DMZ are not needed in the Aruba system. Aruba recommends that organizations consider deploying guests on a nonroutable network with a VLAN that exists only on the Aruba Mobility Controller. Consider having the mobility controller act as the DHCP and NAT server for this self-contained VLAN. The guest role should be locked down so that guest users have limited or preferably no access to internal resources and only limited access to Internet protocols.
- **AP VLANs:** Aruba strongly recommends that edge access VLANs should not be dedicated to APs except in environments where 802.1X is a requirement on the wired edge. The APs should use the existing edge VLANs as long as they have the ability to reach the mobility controller. Deploying the APs in the existing VLANs allows for the full use of the Aruba rogue detection capabilities.

  The one exception to this rule is for AMs. The AMs must be connected to a trunk port that contains all VLANs that appear on any wired access port within range of the AM. This connection is required for the AM to do wireless-to-wired correlation when tracking rogue APs.
- **Quarantine VLANs:** Aruba also recommends against the use of a quarantine VLAN unless it is required by security policy. Instead, Aruba recommends that the integrated firewall and user roles are used to lock down users with a quarantine role. The locations and communications capabilities of the quarantined device are limited more effectively with a quarantine role than with a shared VLAN.

- **Default Router:** In most campus environments, the Aruba Mobility Controller is deployed as a Layer 2 device to provide mobile access and security policy, but not to act as the default gateway for the user subnets. The default gateways typically already exist and are already set in DHCP scopes. To continue to use these devices provides the least disruption to the existing network.
  - Aruba does recommend that the mobility controller act as both default gateway and DHCP server for guest VLANs in all deployments where the VLAN exists only on the mobility controller, and for user VLANs in remote access deployments. In these deployments, the mobility controller is the only networking device with clear visibility into the user subnets, and as such should be deployed as the default gateway.

Mobility

ArubaOS provides seamless wireless connectivity as users move throughout the network through its different client mobility services. The purpose of mobility is to allow seamless roaming across the infrastructure without breaking stateful sessions such as voice and video. These methods allow users to keep the same IP addresses and connectivity as long as their station is active and moving through a contiguous coverage area. Mobility comes in two forms, Layer 2 (VLAN mobility) and Layer 3 (Mobile IP).

Layer 2 Mobility (VLAN Mobility)

Layer 2 mobility is the simplest form of mobility, and it involves a user roaming between mobility controllers that contain the same Layer 2 VLANs. In this case, though the user has transitioned between APs and mobility controllers, the Layer 2 network remains the same and allows them uninterrupted roaming. The downside to Layer 2 mobility is its limited size, because the same VLAN typically only exists over a small area and a limited number of users, even when VLAN pooling is used. VLAN mobility is often deployed only where the number of controllers is relatively small, typically two or less, or a special VLAN is required for nonroutable protocols, typically voice.

Figure 63 *Layer 2 mobility*

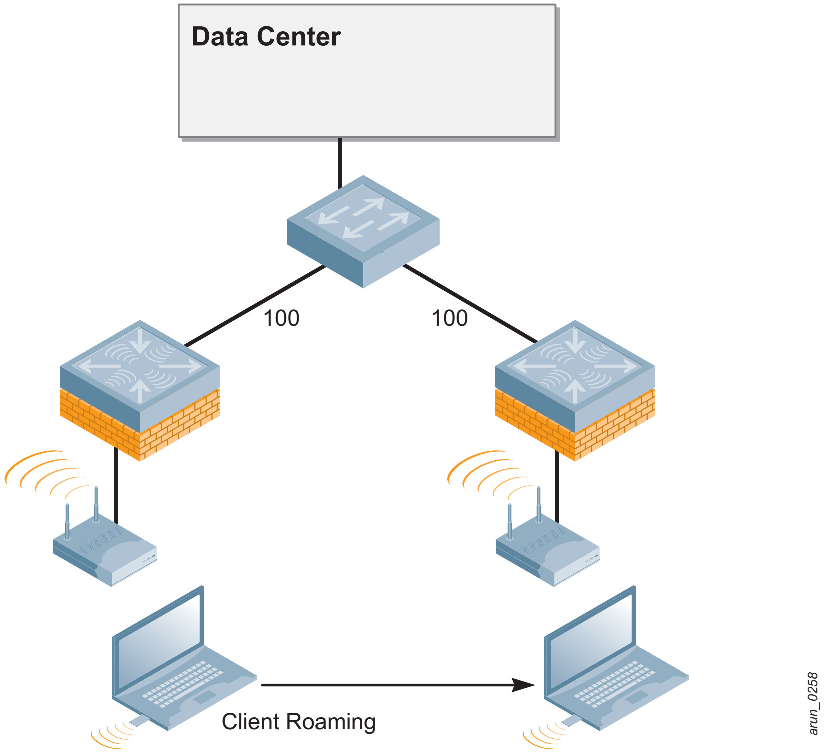

After VLAN mobility is enabled on a virtual AP, any future association to the APs that share that profile causes the mobility controller to look in its bridge table to determine if the device is already assigned a VLAN. If an entry is found, the VLAN membership of the device is preserved, and the BSSID is added to that VLAN. If no bridge entry is found, the station is assigned to the default VLAN of the BSSID or VLAN pool, if any.

The bridge table is aged slowly to preserve the VLAN membership. All controllers that participate in VLAN mobility must be members of the same VLANs (at Layer 2). Use static generic routing encapsulation (GRE) tunnels if there are topological restrictions.

Roaming Between APs in Bridge Mode

Roaming between bridge mode APs is a special case of Layer 2 roaming. In both bridge mode deployment models (campus and remote), it is possible that more than one AP could be deployed in a single location. All of the user session information is stored on the AP itself instead of at the mobility controller, so the users need a way to roam across APs without interrupting their active sessions. To facilitate roaming, Aruba has developed a mechanism for APs to communicate and share user state as the user roams from AP to AP. This mechanism is always enabled when an AP is set to bridge mode, and it requires that all of the APs where roaming will occur be on the same Layer 2 segment.

Figure 64 *Bridge mode Layer 2 mobility*

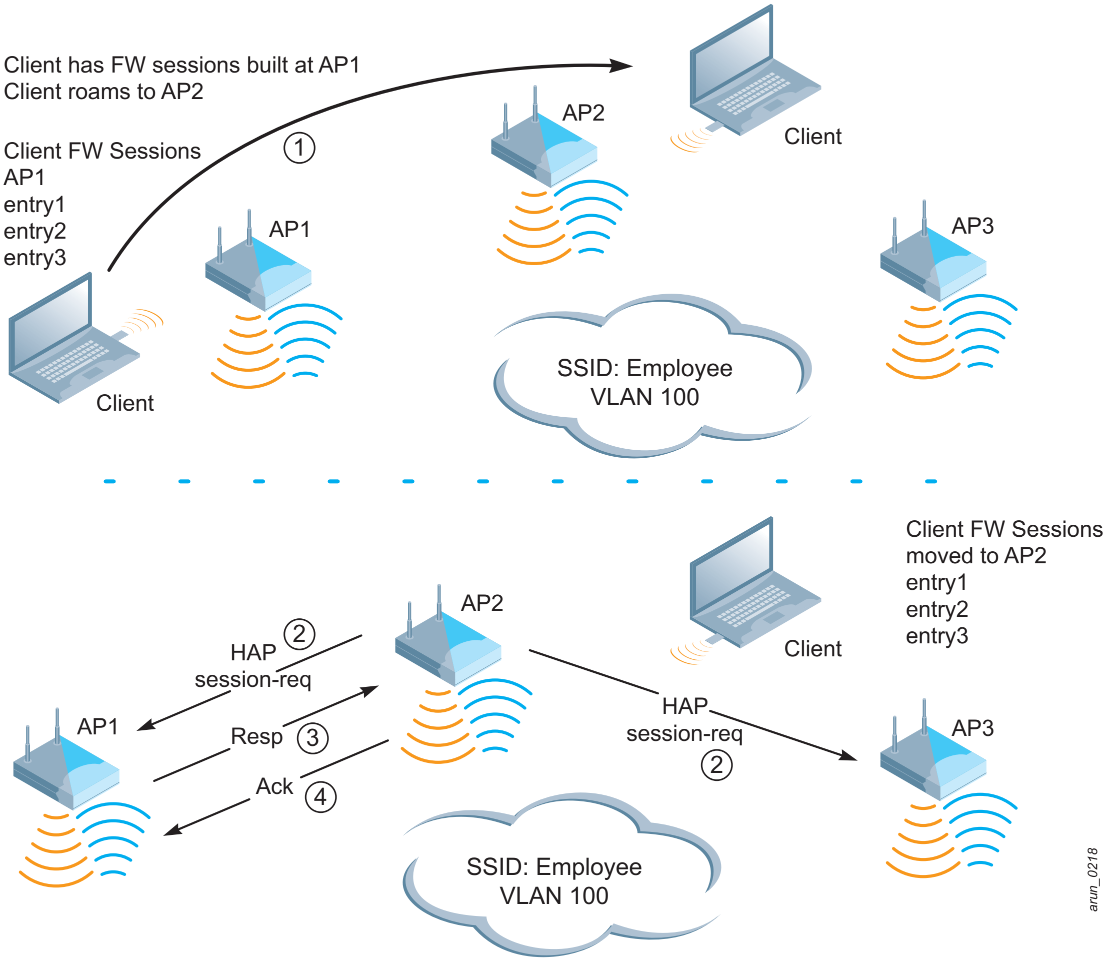

The roaming process occurs as follows:

1. The client begins to roam from AP1 to AP2, which starts an association with AP2.
2. AP2 sends a broadcast message to all APs on the local segment asking if anyone has a current state for the client, identified by its MAC address.
3. Only AP1 responds to the broadcast, and it sends the current session table of the client.
4. AP2 acknowledges the receipt of the session table.
5. AP1 deletes the session state of the client.

6. Roaming is complete.

NOTE

Bridge mode operates independently of the mobility controller, so some functionality is not available because the APs lack visibility into the rest of the network. Please see the *ArubaOS User Guide* for a complete list of features available for bridge mode forwarding. The user guide is available on the Aruba support site at https://support.arubanetworks.com/.

Layer 3 Mobility (Mobile IP)

ArubaOS integrates proxy Mobile IP and proxy DHCP functions, which let users roam between subnets, APs, and mobility controllers without special client software bringing IP mobility to any IP-based Wi-Fi device. The Mobile IP system is the only one that can scale after the system moves beyond a few mobility controllers. After multiple subnets and VLANs exist on different mobility controllers, Layer 3 mobility is required if users must roam between APs that are connected to the different mobility controllers.

With Mobile IP, the ArubaOS automatically tunnels traffic between a roaming client's original mobility controller (the "home agent") and the mobility controller where the user currently terminates ("foreign agent"). With Mobile IP and automatic tunneling, users are able to roam the enterprise without a change of IP address even when they are connected to mobility controllers where their original subnet does not exist.

Mobility Domains

The ArubaOS Mobility Domain is the implementation of Mobile IP addressing that is specified in RFC 3344, also known as Layer 3 roaming. Roaming with a Mobile IP device allows the client to stay connected to services and removes the necessity to reauthenticate Layer 3 services as the point of attachment to the network changes. The Aruba solution extends the RFC functionality in that it requires no special software to be loaded on the wireless client. The Aruba Mobility Controller automatically handles the location changes without client intervention or client-side software configuration.

Figure 65 *Layer 3 mobility*

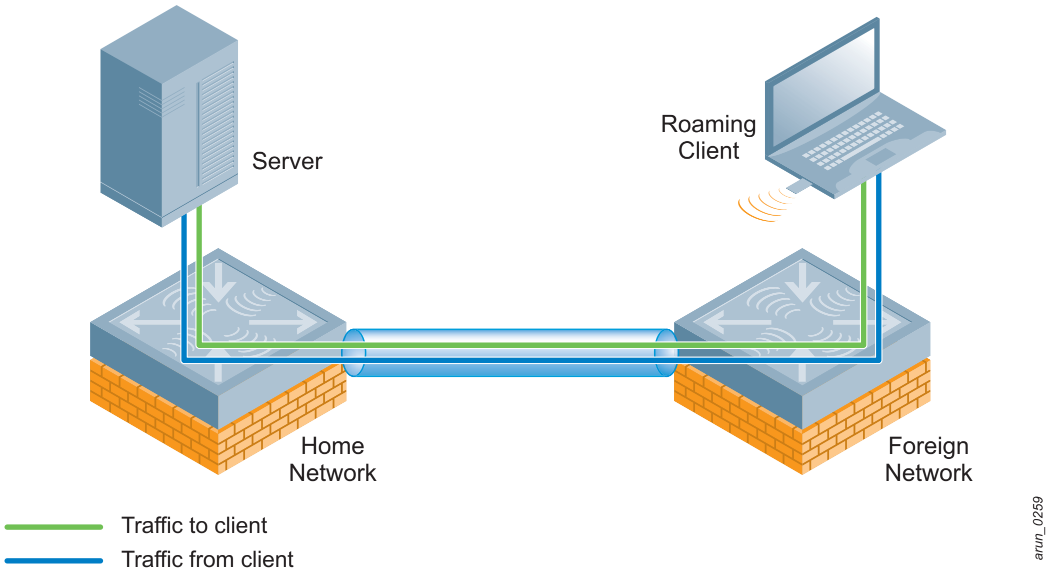

An Aruba Mobility Domain is a logical construct that defines a group of controllers physically close enough to one another that it could be reasonable that a user would roam between them in a single session. You can scale your mobility domain from a single domain on a limited number of mobility controllers to multiple domains. Each controller can handle a separate country, campus, or building depending on your network design and business needs.

Mobility controllers can exist in one or more mobility domains at the same time, much the way a border area router exists in more than one Area in OSPF, but this is not recommended. If a mobility controller can allow a client to roam between two mobility domains, it is better to simply configure the entire network to be a single mobility domain. If a client roams past the last mobility controller that is part of their orginal mobility domain, they will enter a new mobility domain. The client is then forced to receive a new IP address, which breaks any stateful sessions. The mobility domain must be explicitly configured to allow roaming between the various controllers.

Figure 66 *Layer 3 mobility domain*

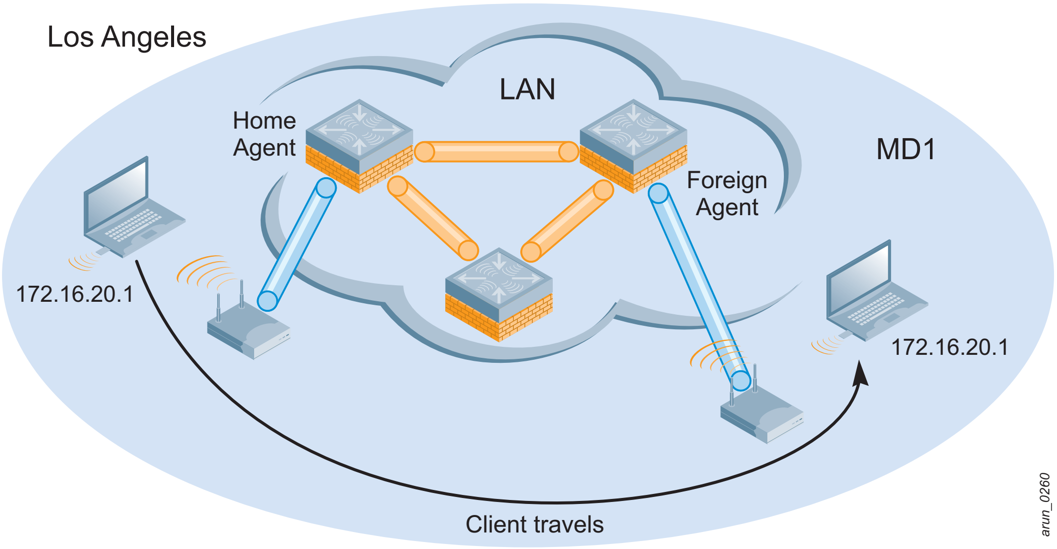

When the client roams off of its “home” network to another network, the client is said to be attached to a “foreign” network. The foreign network is defined as a network managed by a different Aruba Mobility Controller than the one managing the client’s home network, but still within the same mobility domain. The IP address (either VRRP address or controller IP address) of the Aruba Mobility Controller on the foreign network becomes the client’s “care-of address.” This address is passed to the Aruba Mobility Controller on the home network, where the home agent keeps a map of clients and care-of addresses. The home agent learns the care-of address from a similar process on the foreign network known as the foreign agent.

From an IP perspective, the client still appears to be attached to its home network, so all data bound for that client is routed to its home agent. When the home agent sees packets bound for the client, it tunnels those packets in an IP-IP tunnel to the foreign agent for delivery to the client. Any traffic generated by the client is sent directly from the foreign agent using standard IP routing and delivery mechanisms. Routing tables remain intact, and the client can continue to use the IP address that it acquired in its home network.

Mobility domains take some amount of planning, but they generally follow the physical layout of the network. For a centralized network that is located in a single building or campus, it may be possible to design a network that has only a single mobility domain. The main design consideration should always be "can the user realistically expect to roam between the subnets and controllers in a single session?" This roaming is possible in the same building or on a campus with coverage between buildings; however, roaming between an office in Los Angeles and an office in New York is not possible.

Figure 67 *Separate mobility domains for geographically dispersed sites*

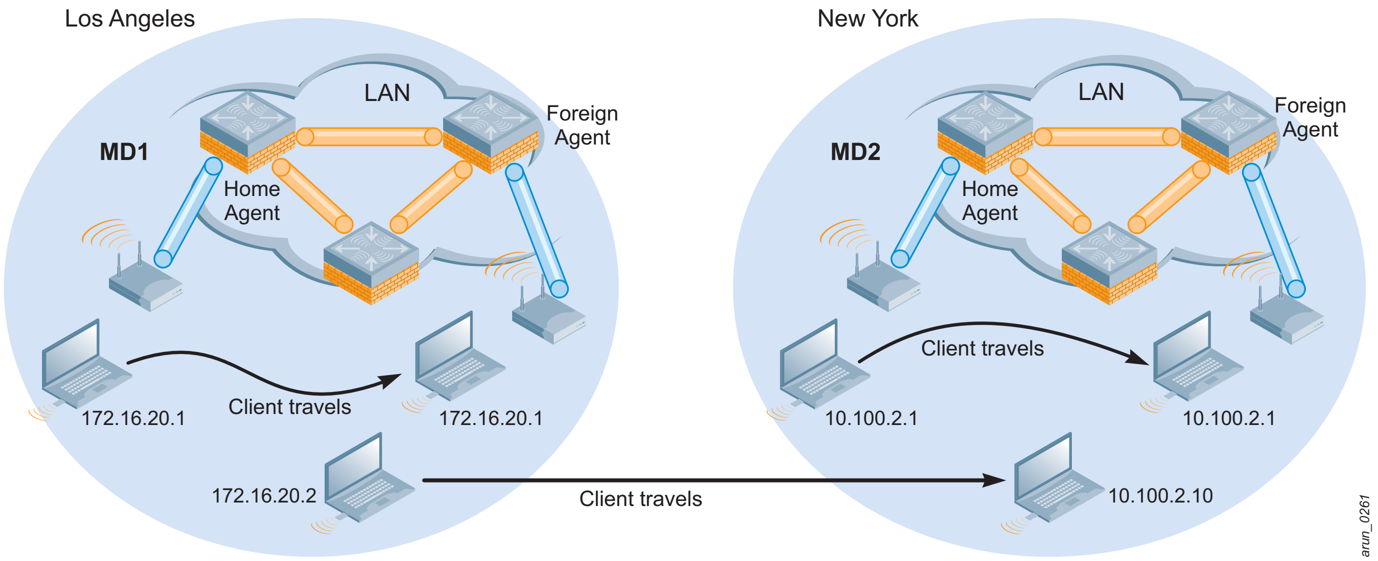

To plan a mobility domain, begin by looking at the physical topology of the network and the network map, and focus on the APs and controllers. Generally, these elements provide the information you need to develop a logical grouping of mobility domains. Where outdoor coverage exists between buildings on a campus, it is recommended that the heat maps of the network are also examined. Determine if the coverage areas provide enough connectivity and overlap to allow your clients to transition buildings. Outdoor APs may extend this coverage between buildings and provide a larger mobility domain.

Mobility Recommendations

Aruba recommends Layer 3 mobility any time more than two mobility controllers are present in the network. While this configuration creates more overhead for the administrator initially, it leads to a cleaner network design where VLANs exist in fewer places and are less likely to be overloaded.

For one or two mobility controller networks, the recommendation is to use Layer 2 mobility. In these deployments, VLANs need not span large numbers of controllers, so nothing is gained from configuring Layer 3 mobility. All user VLANs must exist on all mobility controllers, and typically these mobility controllers are configured to be redundant.

Chapter 10

Redundancy Models

As WLAN moves from a convenience network to a mission-critical application, the need for availability and redundancy also increases. Aruba provides several redundancy models for local and master mobility controllers. Each of these options, including the choice to forgo redundancy, must be understood so that the correct choice can be made for each deployment model.

Redundancy is always a tradeoff between the cost of building a redundant network and the risk of the network being unavailable if an outage occurs. For each organization, the tradeoffs must be weighed between the cost of implementing a redundancy solution and the cost of an outage. In some cases, multiple types of redundancy are possible, and it is up to the organization to gauge its tolerance for risk given the pros and cons of each redundancy model.

Figure 68 *Scale of redundancy for mobility controllers*

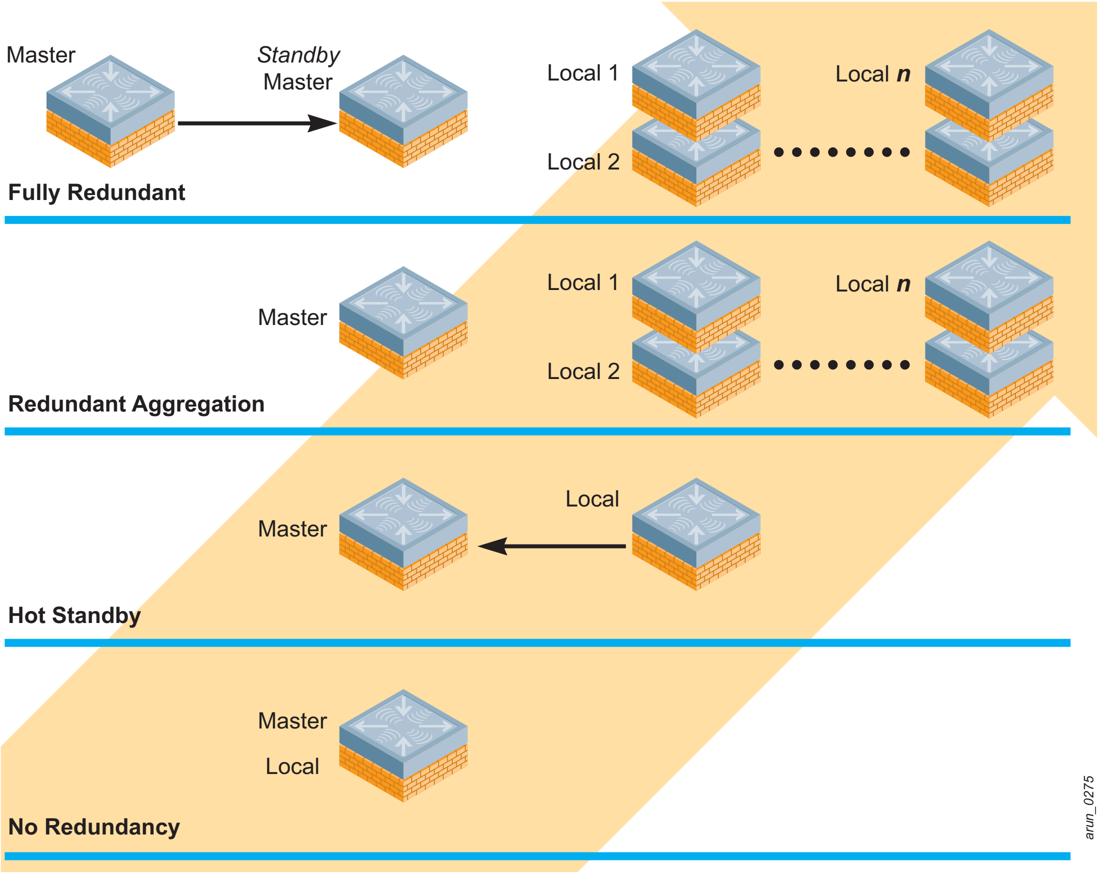

The scale of redundancy has different levels:

- Having no redundancy at all
- Adding redundancy between a set of mobility controllers
- Adding redundancy for aggregation level mobility controllers
- Having a completely redundant network

At each level, the cost and complexity increases, and the chance of the network being unusable due to a network outage decreases. The following sections discuss redundancy at each level and what the consequences are of running a network without redundancy.

Master Redundancy

The master mobility controller is the center of the control plane. The master controller handles initial AP boot up in Layer 3 deployments, policy configuration and push to the local mobility controllers, local database access, and services such as WIP coordination and location. Additionally, if CPsec is enabled on the network, the master is responsible for certificate generation.

Figure 69 *Master redundancy using VRRP and database synchronization*

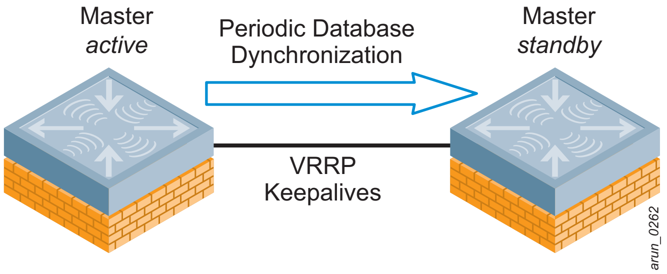

To achieve high availability of the master mobility controller, use the master redundancy method. In this scenario, two controllers are used at the management layer, and one controller configured as an active master and one configured as a standby master. The two master mobility controllers operate in a hot standby redundancy model. One master is the active primary, and the second is a standby that receives updates from the master about the state of the network.

- The two masters synchronize databases and run a VRRP instance between them. The virtual IP (VIP) address configured in the VRRP instance is also used to communicate with the current primary master. This address is given to the local mobility controllers, wired APs, and wireless APs that attempt to discover a mobility controller. The VIP is also used for network administration.

Figure 70 *Master redundancy failure scenario for the local mobility controller*

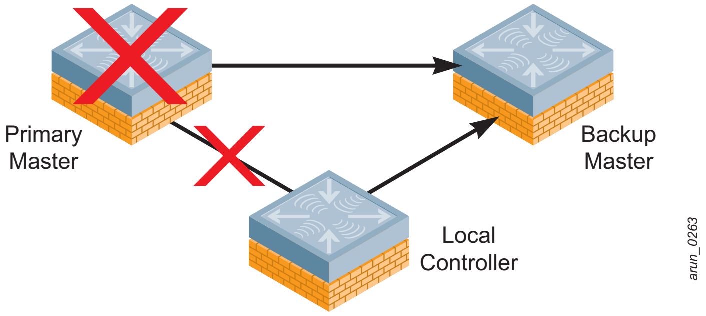

When the primary master becomes unreachable for the timeout period, the backup master promotes itself to be the primary master and uses the VRRP IP address. All traffic from local mobility controllers and APs to the master automatically switches to the new primary.

Aruba does not recommend enabling preemption on the master redundancy model. If preemption is enabled and a failover occurs, the new primary remains the primary even when the original master comes back online. The new primary does not revert to a backup unless it is forced to by an administrator.

Local Redundancy

Three types of local mobility controller redundancy are available. Each type of local redundancy is appropriate in a particular scenario, and sometimes they operate together.

Active-Active (1:1)

In the Aruba active-active redundancy model, two local mobility controllers share a set of APs, divide the load, and act as a backup for the other mobility controller. When two controllers operate together, they must run two instances of VRRP and each controller acts as the primary for one instance and backup for the other.

Figure 71 *Active-active redundancy, both mobility controllers reachable*

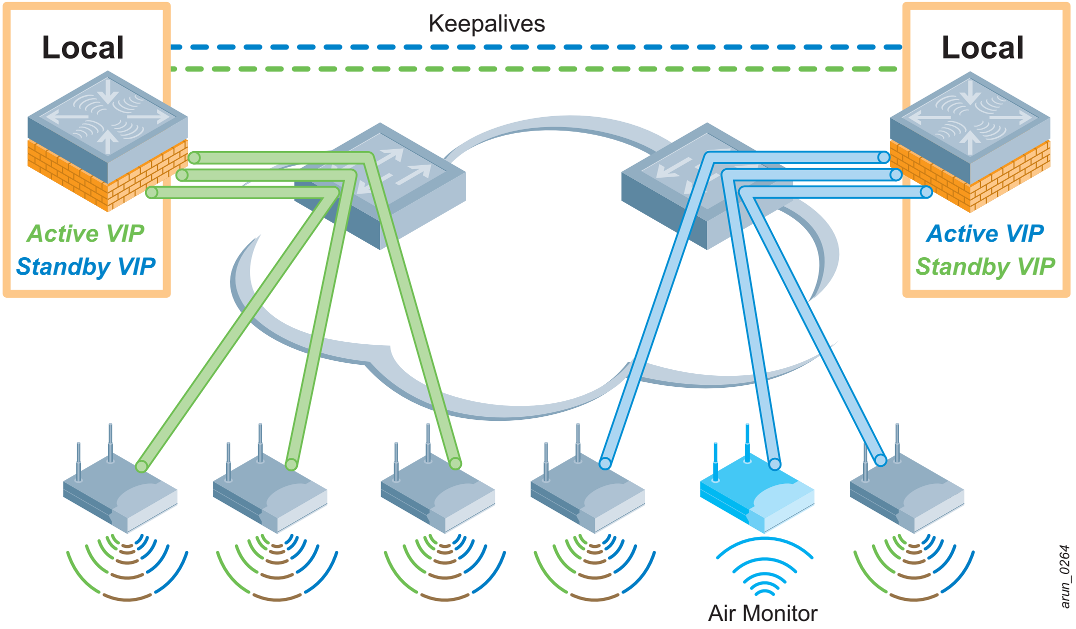

Using this model, two local controllers terminate APs on two separate VRRP VIP addresses. Each Aruba Mobility Controller is the active local controller for one VIP address and the standby local controller for the other VIP. The controllers each terminate half of the APs in this redundancy group. The APs are configured in two different AP groups, each with a different VIP as the local management switch (LMS) IP address for that AP group.

Figure 72 *Active-active redundancy, mobility controller unreachable*

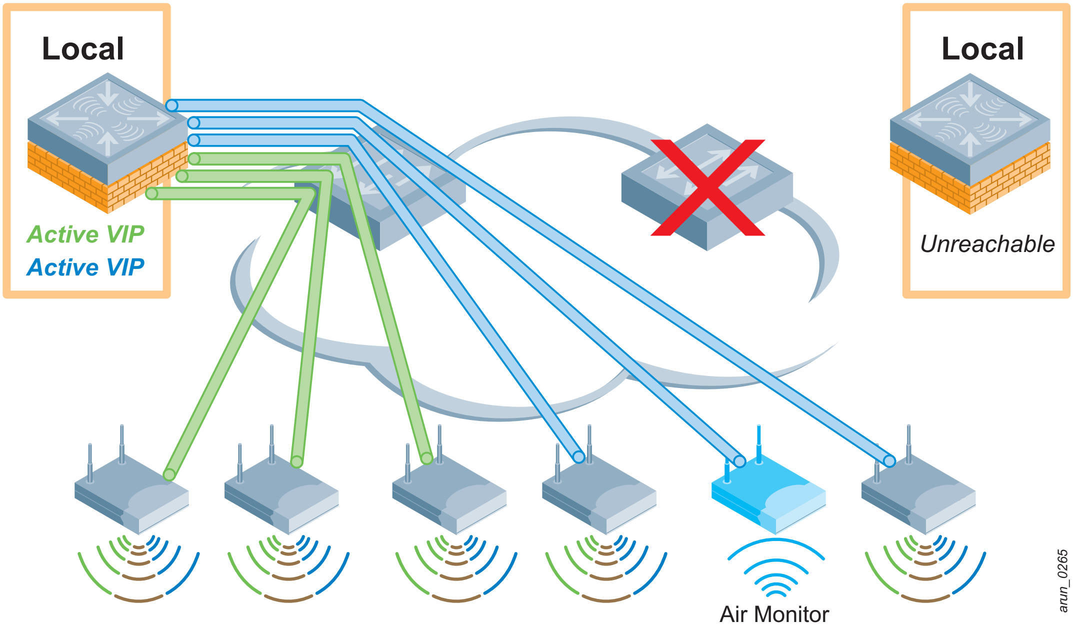

When one active local controller becomes unreachable, APs connected to the unreachable controller fail over to the standby local mobility controller. That controller now terminates all of the APs in the redundancy group. Therefore each controller must have sufficient processing power and licenses to accommodate all of the APs served by the entire cluster.

In this model, preemption should be disabled so that APs are not to forced to fail back to the original primary when it comes back online. APs will not fail back, so this model requires that the mobility controller be sized appropriately to carry the entire planned failover AP capacity for an extended period of time. See the following note for a discussion of appropriate scaling.

NOTE

When determining the AP load for active-active, some thought should be given (from a capacity standpoint) to what will happen to the backup controller when the APs fail over. If each mobility controller is at 50% of total capacity, when a failure occurs, the mobility controller that the APs fail over to will now be at 100% capacity. As with any system component, it is never a good idea to run the system at maximum capacity. Aruba recommends that each mobility controller be run at 40% capacity, so that when a failover occurs, the surviving mobility controller will only be at an 80% load. This load gives the mobility controller the room to operate under the failover conditions for a longer period of time. An 80% load also reduces the time for APs to fail over from the primary mobility controller to the backup mobility controller.

Active-Standby (1+1)

The active-standby model also has two controllers, but in this case one controller sits idle while the primary controller supports the full load of APs and users.

Figure 73 *Active-standby redundancy, primary mobility controller reachable*

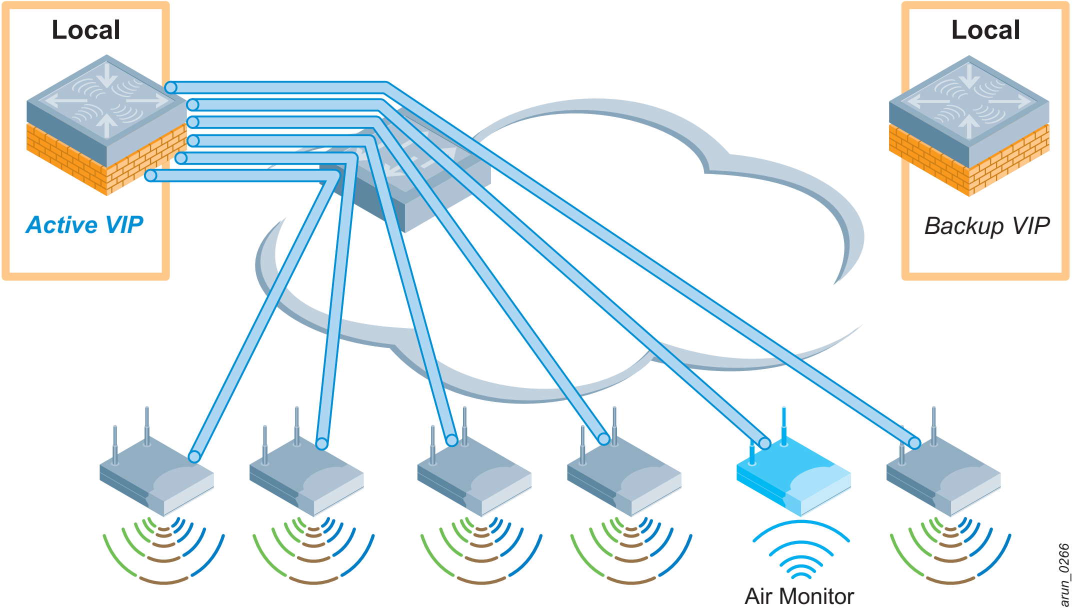

When a failure occurs in the active-standby model, all of the APs and users must fail over to the backup controller. This model has a larger failure domain and will have some increased latency as the full load of APs fails over and users reauthenticate. This form of redundancy uses the LMS and backup LMS configuration for the AP. Alternatively, a single VRRP instance could be run between the two controllers, and all APs for the pair would terminate against this VRRP IP address.

Figure 74 *Active-standby controller, primary mobility controller unreachable*

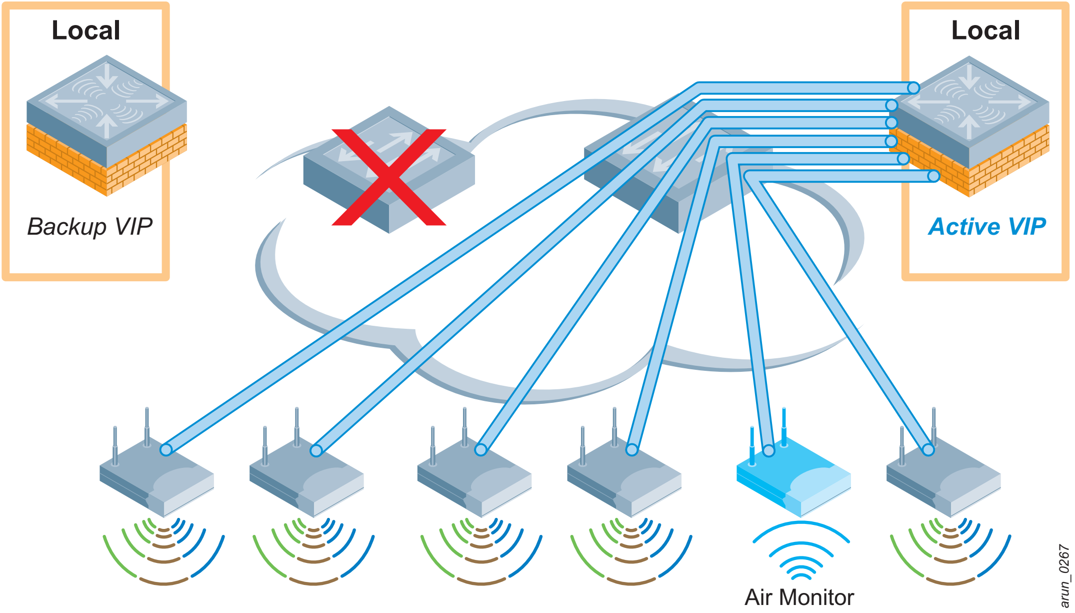

The active-standby model is primarily used when the two mobility controllers are separated by a Layer 3 boundary, which makes it impossible to run VRRP, which operates at Layer 2, between the two mobility controllers. Mobility controllers are typically separated by a Layer 3 boundary when they are deployed in separate data centers.

As with active-active, when the active local mobility controller becomes unreachable, all of the APs that are connected to the unreachable controller fail over to the standby local mobility controller. That controller carries the full AP load of both mobility controllers for the duration of the outage. Therefore each controller must have sufficient processing power and licenses to accommodate all of the APs served by the entire cluster.

Many-to-One (N+1)

The many-to-one model is typically used in remote networks where branch offices have local mobility controllers, but redundancy on site is not feasible. It is possible to use N+1 on the campus as well, but here consideration should be given to the ratio and likelihood that sections of the campus might become unreachable, which would cause a multiple controller failover. This scenario requires that a secure connection be established between the sites that is independent of the local mobility controller, and that the connection should have high bandwidth and low latency.

Figure 75 *N+1 redundancy, local mobility controller active*

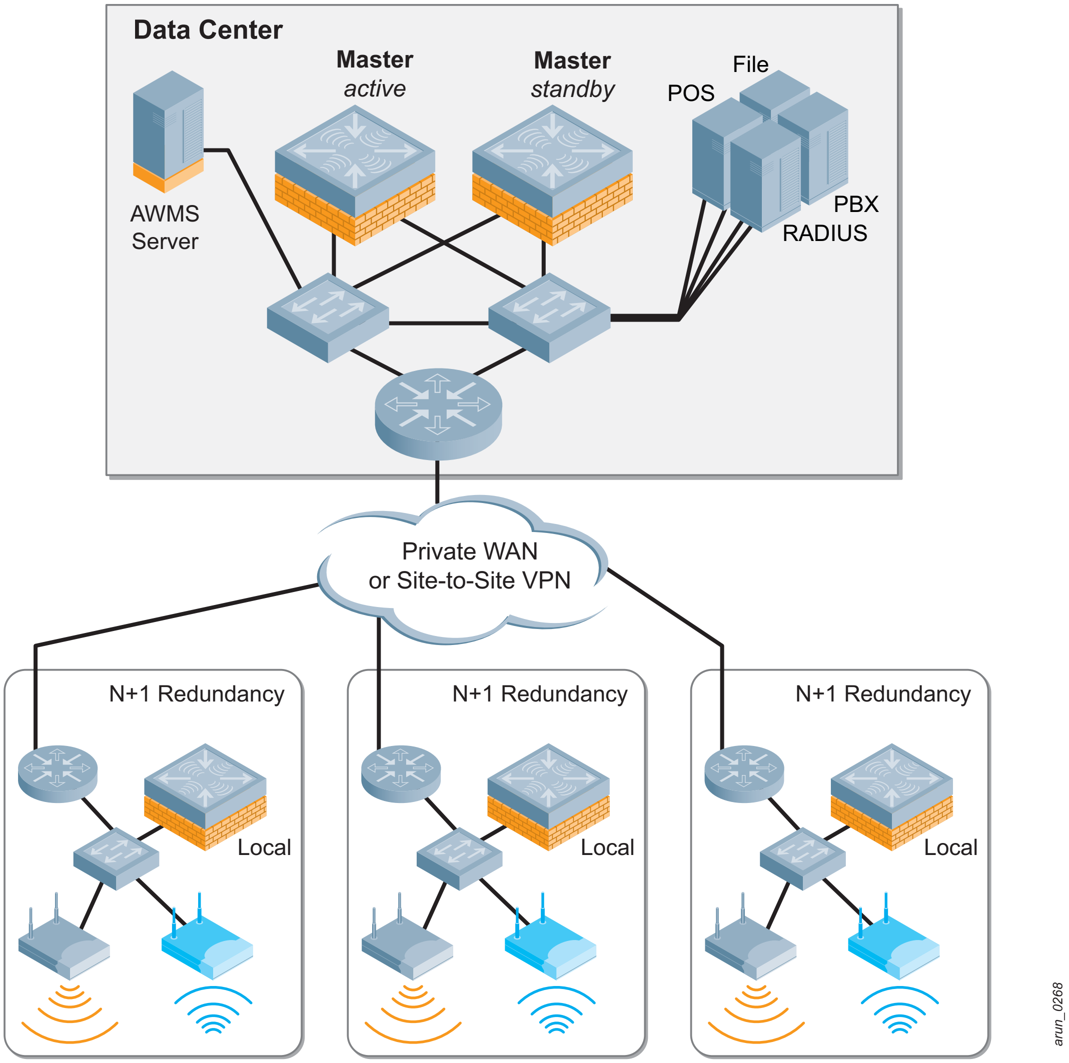

When the local mobility controller at the remote site fails, the APs fails back to the backup LMS configured for that purpose, just as in the active-standby scenario.

Figure 76 *N+1 redundancy, local mobility controller failed, AP connected across the WAN*

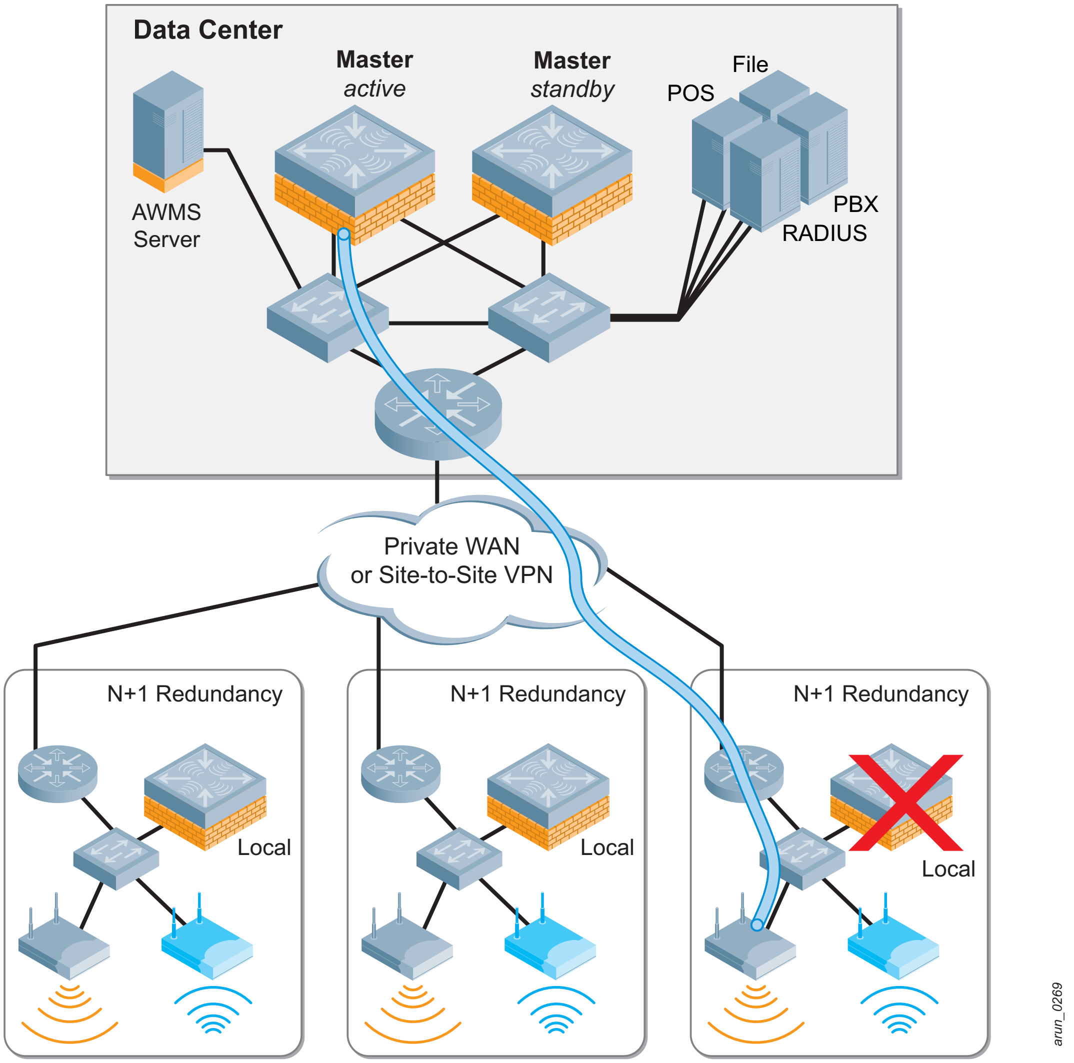

The difference in the N+1 scenario is that this failure is typically across a WAN link, and the backup controller should be large enough to handle multiple site failures at the same time. Though a typical small site might have a handful of APs on a smaller mobility controller, the central site must have a much larger mobility controller with increased licensing to handle the expected number of failures of locals. In typical designs, only a single failure is anticipated, but some organizations require more resiliency against failure of multiple sites. Common cases include retail stores, where more than a single store may have an outage at any one time due to the sheer number of sites and the fact that the controller may be in user-accessible space.

Aruba strongly recommends that preemption be enabled in this scenario. Due to the limited capacity of the redundant mobility controller and the possible delay introduced by failing over to a remote site, it is recommended that APs be moved back to their original mobility controller as soon as service is restored.

Comparison of Local Redundancy Models

The following table summarizes the pros and cons of each redundancy model, which allows the network manager to make the proper redundancy decision for their network.

Table 19 *Comparison of Redundancy Models*

Redundancy Type	Pro	Con
Active-Active (1:1)	● Smaller failure domain, because fewer APs must fail over in the event of an outage ● Outage duration is smaller, because fewer APs will take less time to recover, typically about half as long as failing over a fully loaded mobility controller ● All mobility controllers in use at all times ● Reduced load on each mobility controller	● More expensive, because all mobility controllers must be licensed to handle the full compliment of APs in the failure domain. Aruba recommends that this load be planned to 80% of each mobility controller’s maximum capacity ● Twice as many mobility controllers are required vs. no redundancy
Active-Standby (1:1)	● If APs fail to the backup controller, essentially nothing has changed in the network except where the APs and users are hosted	● Has the same cost structure as the active-active redundancy model, with two sets of mobility controllers and two sets of licenses ● Larger failure domain, all APs must fail to the backup mobility controller, typically takes twice as long as active-active ● Outage duration will be longer, because more APs must be recovered
Many-to-One (N+1)	● Cost-optimized model, fewer redundant mobility controllers are required, and need only be licensed and scaled to handle the maximum number of failed mobility controllers ● Typically only one redundant mobility controller is deployed	● Multiple failures can overwhelm the redundant mobility controller, which causes a network down scenario ● Preemption must be enabled to clear APs back to the primary mobility controller as soon as it is recovered, which results in a second unplanned outage

Aruba recommends using active-active redundancy wherever possible. Active-active provides the fastest recovery time in the event of a network outage with the least disruption to the end user. Aruba also recommends in all models that mobility controllers not be loaded past the 80% mark to help increase stability of the network.

AirWave Redundancy

The AWMS sits above the mobility controllers in the network hierarchy and provides the management of the network. When AWMS is used as the monitoring and configuration system for the network, redundancy should also be considered. AWMS servers are made redundant through a special license installed on one or more servers. These redundant servers then monitor a set of active AWMS servers by polling to check that the servers are up and offloading the daily backup file.

Figure 77 *AWMS servers and AWMS backup server, with Master Console*

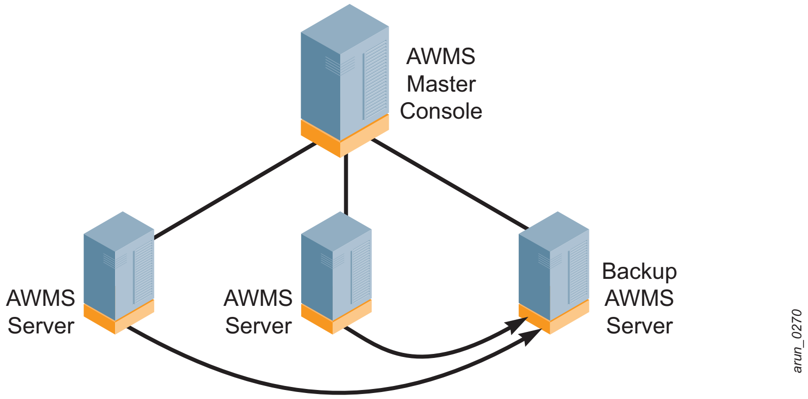

If an active AWMS system becomes unreachable by the redundant system, the redundant system restores the last backup for that active server and begins polling the devices that were previously monitored by that server. No automatic failback occurs if the original AWMS server recovers. Network administrators must manually restore the server if they decide to restore the original AWMS server to production.

Data Center Redundancy

If the data center of an organization experiences an outage, where all mobility controllers at a particular site are offline but the network continues to operate, the APs can, in theory, fail over to a redundant set of mobility controllers in another location. The redundant controllers can be either in the same data center but connected by discrete power and data connections, or in a remote data center that is reachable by a private WAN or IPsec link.

Figure 78 *Active-active plus LMS and standby backup LMS*

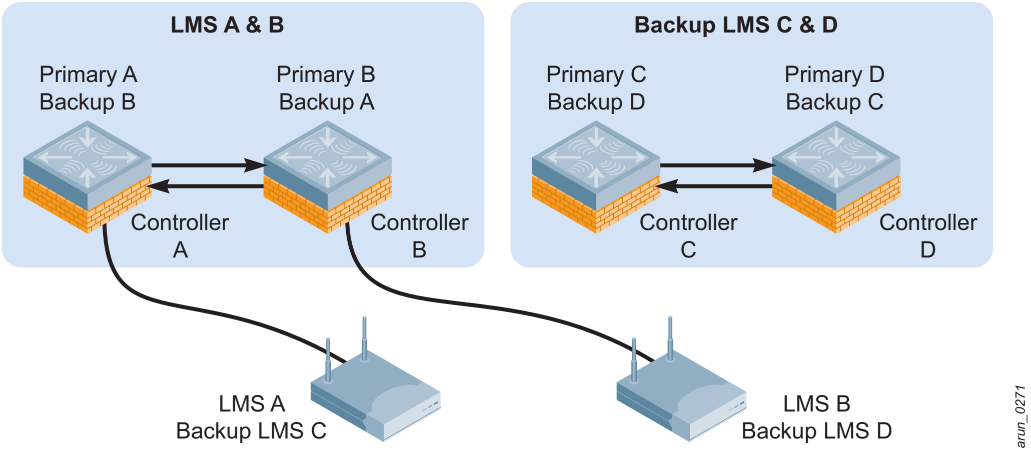

Figure 79 *Two active-active pairs, with LMS and backup LMS all serving APs*

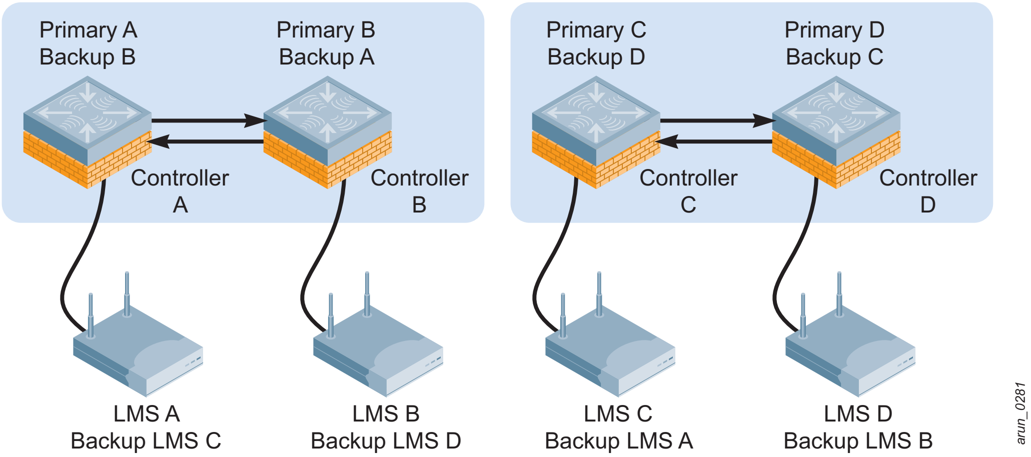

Data center redundancy consists of four total controllers and four instances of VRRP. The APs can be set up either to split between two of the mobility controllers (active-active) with a pair in hot standby, or spread evenly across all four mobility controllers. In this model, the APs are set up so that they operate on the VIP of their primary pair of mobility controllers, and their backup is one of the two VIPs on the second pair of mobility controllers.

Figure 80 *Failure of single primary mobility controller in active-active with LMS and backup LMS*

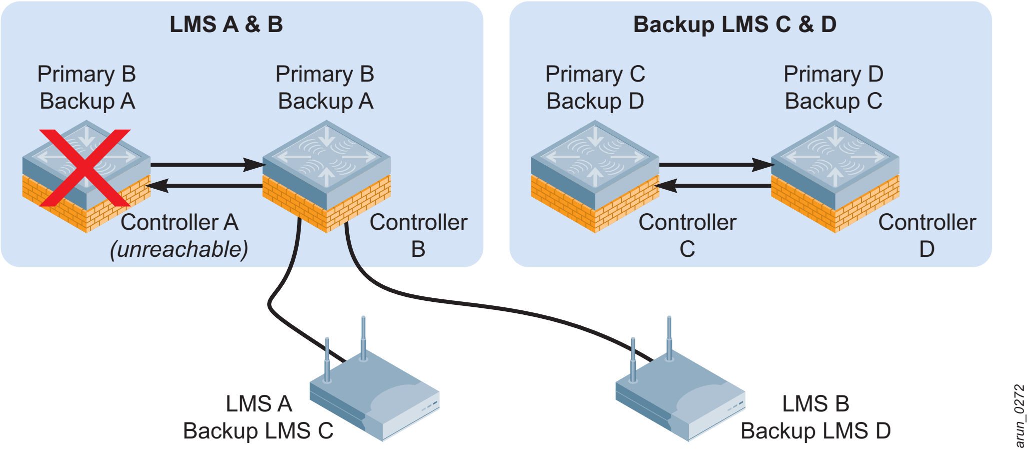

Figure 81 *Failure of primary the primary data center in active-active with LMS and backup LMS*

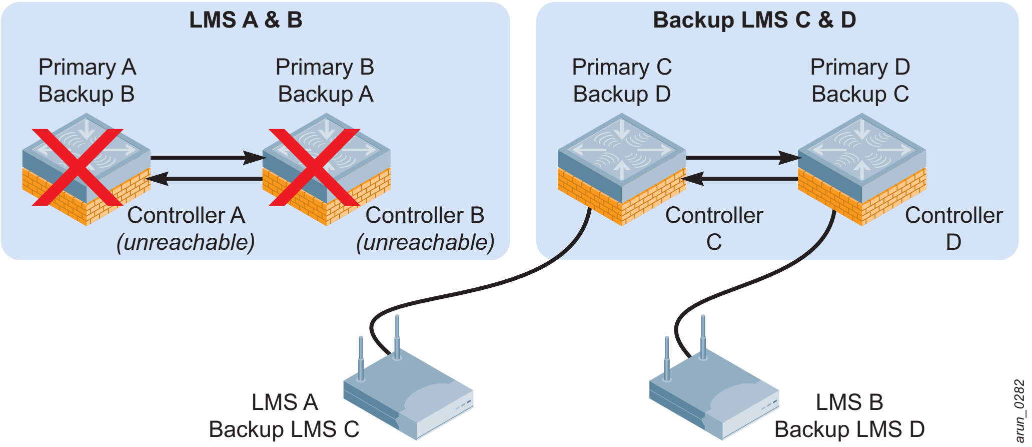

In a failure scenario, the failure of one mobility controller in a pair results in typical active-active failover. If the second mobility controller in the pair fails, the APs fail over to their backup pair of controllers and split between the two VIP instances. In either deployment model, all four mobility controllers must be licensed and capable of supporting the full AP load.

Figure 82 *Failure series, active-active with LMS and standby backup LMS*

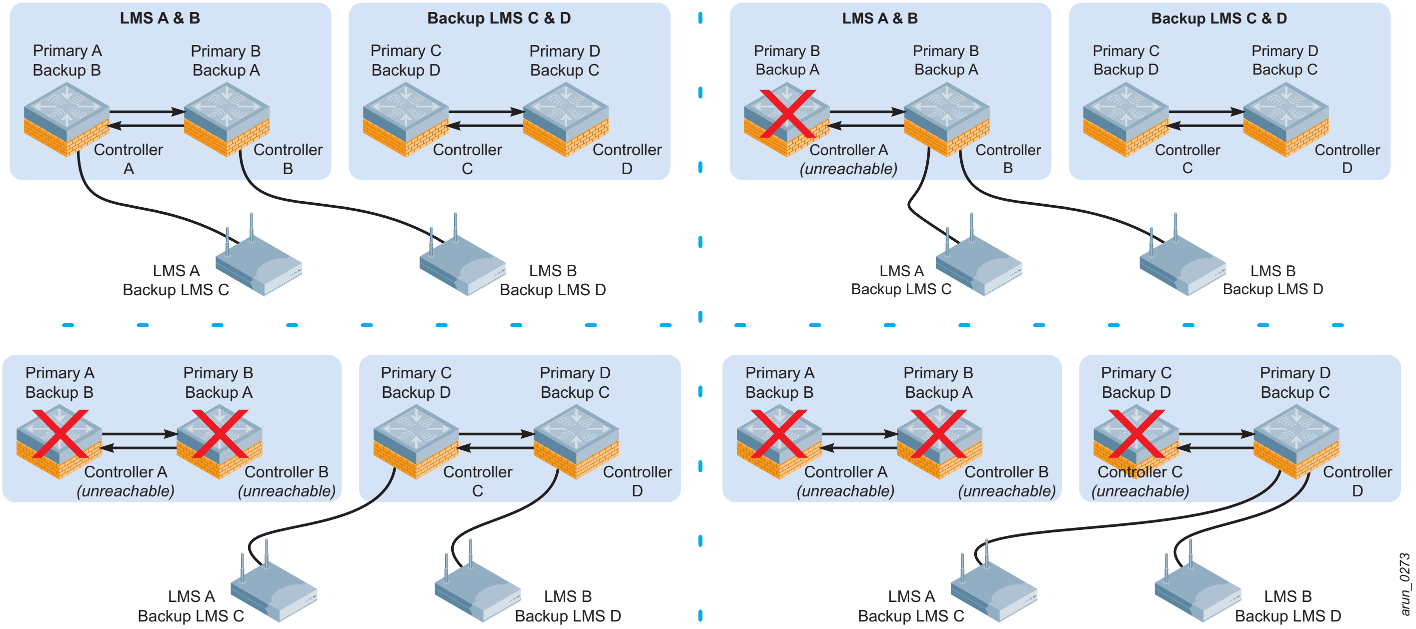

Scenario 1:

1. 412 APs were split across two active M3 mobility controllers (206 APs each), with each group active on one of the two VRRP instances in the first pair of mobility controllers and the second pair standing by to receive APs.
2. In the first failure (A), the APs on the failed mobility controller fail over to the active backup mobility controller (B). This change results in 412 APs on the backup mobility controller, and 0 APs on each mobility controller in the second cluster.
3. When the second mobility controller in the group (B) fails, the 412 APs from the failed cluster distribute themselves evenly across the two mobility controllers that are still active in the second cluster (C & D). This change results in 206 APs on each mobility controller.
4. If a third mobility controller fails (C), all 412 APs become active on the remaining mobility controller (D).
5. As a result, each mobility controller must be licensed to support all 412 APs if three of the other mobility controllers become unreachable.

Figure 83 *Failure series, active-active with LMS and backup LMS also in use*

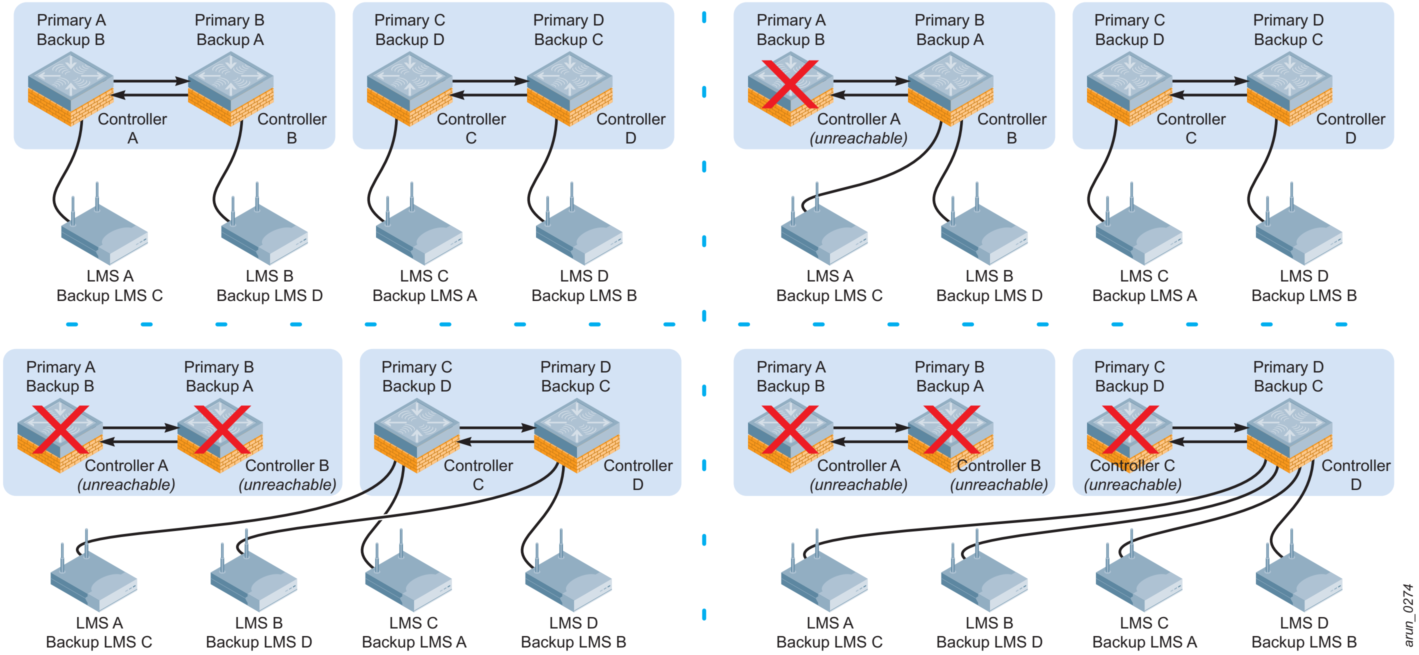

Scenario 2:

1. 412 APs were split across four active M3 mobility controllers (103 APs each), and each group was active on one VRRP.
2. In the first failure, the APs on the failed mobility controller (A) fail over to the active backup mobility controller. This change results in 206 APs on the backup mobility controller (B), and 103 APs on each mobility controller (C & D) in the second cluster.
3. When the second mobility controller (B) in the group fails, the 206 APs from the failed cluster distribute themselves evenly across the two mobility controllers (C & D) that are still active in the second cluster. This change results in 206 APs on each mobility controller.
4. If a third mobility controller (C) fails, all 412 APs become active on the remaining mobility controller (D).
5. As a result, each mobility controller must be licensed to support all 412 APs if three of the other mobility controllers become unreachable.

No Redundancy

In early WLAN deployments, redundancy was often viewed as a luxury, because the network was not deemed to be mission critical. Not having redundancy is still considered acceptable to some organizations, though it is not recommended by Aruba. This section describes what is lost when a component of the system fails without redundancy enabled.

Master Mobility Controller – No Redundancy

If the master mobility controller fails without a backup, the following services stop working:

- **AP boot:** During the AP boot cycle, the AP must discover and connect to a provisioning mobility controller. In almost all deployments this is the master mobility controller, because that mobility controller typically is not serving APs and is able to be a single source for AP provisioning. It is also far easier to configure either DNS lookup or a single DHCP option to find a single mobility controller than to manage multiple lookups or scopes. It is also possible to use Layer 2 discovery mechanisms to find a local mobility controller, but this is not realistic in larger deployments. If the master is unreachable, in certain cases the APs may not be able to reboot until the master is restored or their boot process is modified:

  - In situations where DHCP option 43 is used, the APs will not be able to boot until a new master is in place or the DHCP scope option is modified to point at either a new master mobility controller, or to a local mobility controller.
  - If DNS is used to locate the master, the APs will be down unless a second IP is also returned in the DNS response that points to a local mobility controller. Note that this configuration results in a protracted outage, and the local mobility controller must have AP capacity to bring up and then redirect the APs as they fail to the backup DNS response. This outage is longer in duration, with each AP taking approximately 4-5 minutes to fail to the backup. Depending on the AP capacity on the backup, several attempts may be needed before the AP is able to connect and be properly redirected.
  - APs that are relying on Aruba Discovery Protocol (ADP) will continue to operate as long as a local mobility controller is capable of answering their ADP request. These APs require Layer 2 connectivity to the local mobility controller for ADP to function.
  - In all cases, APs that are currently operating will continue to do so in the event that the master becomes unreachable until they are rebooted or power cycled.

- **Local mobility controller policy configuration:** Configuration, done either on the master or AirWave, requires that the master is operational to push configurations to the local mobility controllers. If the master is not available, changes to the network policy configuration will not be possible without manually modifying each mobility controller separately, though local configuration at the IP level will still be possible.
- **Local database access is lost:** If the master becomes unreachable, guest access using the local database, as well as when roaming between locals when machine authentication is enabled, will be lost.
- **Monitoring, heat maps, and location:** If AirWave is not present in the network, centralized network monitoring, heat map generation, and location services will all be down.
- **Valid AP table:** When the master is down, the valid AP table will no longer be available for updates. The local mobility controllers will continue to function with cached data until that ages out, at which time other APs in the network will be seen as “unknown” instead of valid, interfering, rogue, etc. When this occurs, adaptive radio management (ARM) increases power to the edge APs on both sides in an attempt to increase coverage and work around the now unknown AP. At AP border areas, overlapping channels and power will lead to increased interference.
- **WIP coordination:** When the master is down, wireless intrusion prevention (WIP) will lose its coordination capabilities between locals. Any new APs showing up will be classified as “unknown”, which prevents automatic containment from functioning. Existing data will still remain until it ages out, at which point all of the APs will begin to be reclassified as “unknown.” If protection of valid stations is enabled, clients will be prevented from joining any AP that is not valid, which after some time will be all APs that a mobility controller can see that are not directly attached.
- **AP white lists:** The two varieties of white lists are the CAP and Remote AP (RAP) white lists. For the CAP white list, all mobility controllers share a copy of the white list, but without the master, they lose the capability to synchronize the lists. The RAP white list must be manually exported to the local mobility controller to ensure that operations continue, but no additional APs can be authorized while the master is unreachable.
- **CPsec:** Failure to have a backup for CPsec will result in the same failures as a master mobility controller, with the additional problem. If the master physically must be replaced, as soon as it is brought online, the entire network will go back through the recertification list. In addition, the AP white list will need to be rebuilt.

Local Mobility Controller – No Redundancy

If a local mobility controller becomes unreachable and has no backups configured for the APs, all APs assigned to that mobility controller will be down and no users will be able to connect. Any AMs associated to the controller will also be down, which eliminates the capability to scan for threats and contain rogue devices. This situation will continue until the APs are reprovisioned and assigned to another mobility controller or the original or replacement local becomes reachable again.

AWMS – No Redundancy

If an AWMS appliance becomes unreachable without a failover server in place, no monitoring, logging, or reporting will be available for any of the mobility controllers, APs, or users that the AWMS was monitoring. In addition, other functionality may be lost depending on the configuration of the AWMS system:

- **Configuration:** If the AWMS system is the configuration point for the network, policy-based configuration will be down for the duration of the outage. Configuration is still possible on the master mobility controller, but when AWMS is restored, the configuration files will be out of sync.
- **WMS Offload:** If the WLAN Management Suite (WMS) database offload is enabled to move the WMS database from the master to the AWMS server, the valid AP list will be lost in the same manner as if the master had failed without a backup. All controllers will keep the valid list for some period until it ages out, at which point the effects described under the master mobility controller section for ARM and WIP will begin to take effect.

Data Center – No Redundancy

Data center redundancy is commonly only deployed by organizations with extremely high availability requirements and the ability to have the APs connect through a separate set of infrastructure to the second set of controllers. Each organization must make a decision about the acceptable level of risk vs. cost around this higher level of redundancy.

Aruba Recommendations for Redundancy

Wireless networks are no longer convenience networks. They are now mission-critical components of the network. As such, they need to be treated like any other mission-critical system. Aruba recommends redundancy at all levels of the system to ensure a highly available network for users.

Table 20 *Redundancy Recommendations*

Controller	Campus	Branch Office	Remote Access (DMZ)	Data Center
Master	Master redundancy	N/A	Master redundancy	Master redundancy
Local	Active-active redundancy, each mobility controller loaded at 40% of capacity, licensed to 80% of capacity	Active-active redundancy where possible, N+1 redundancy minimum	Active-active redundancy, each mobility controller loaded at 40% of capacity, licensed to 80% of capacity	Active-active redundancy, each mobility controller loaded at 40% of capacity, licensed to 80% of capacity
AirWave	At least one redundant server per network	N/A	N/A	At least one redundant server per network

Chapter 11

Mobility Controller Selection

The selection of the proper mobility controller depends greatly on the application and usage model for the network. This section examines the network usage considerations that must be taken into account, controller scaling, and selection criteria.

Information Gathering

Selecting the proper mobility controller for the deployment depends on a number of factors, including forwarding mode, usage model, and AP count. Take these factors into account to select the proper mobility controller for the application.

- **AP count:** The most common selection criteria for many organizations, the number of APs, often dictates a minimum controller scale. To determine the required AP count, a planning tool such as the Aruba VisualRF Plan or a traditional site survey will determine the number of CAPs and AMs necessary to provide adequate coverage. For remote access solutions, use the number of RAPs and/or the number of VIA agent or VPN users. Mixed environments require additional planning to ensure that the combined CAP and RAP counts do not exceed the maximum supported limits on the mobility controller.
- **User count:** Each platform has a maximum user count that limits the maximum users that can associate with each controller. Look at the number of users that will use the WLAN at each site. Include employees, guests, contractors, smart phones, and autonomous systems. For most organizations, an individual person should count as two devices in the user count calculation. This count can be increased if additional IP-based devices, such as wireless desk phones and tablets, are also in use.
- **VIA users:** If the VIA agent will be deployed, the number of users must be known so that the deployment can be scaled appropriately. Additionally, the decision to make SSL fallback available for the VIA agents has an impact on the system and must be taken into account when selecting the appropriate mobility controller model and quantity.
- **Forwarding modes and CPsec:** The forwarding mode selected for the mobility controller affect how much traffic and how many tunnels the AP will generate. In addition, CPsec processing adds additional processing overhead during boot up and when APs are being certified. Remember that CPsec is required for some modes of operation.
- **Data throughput:** As with any networking device, each mobility controller has a maximum platform throughput, which is also affected by encryption and firewall processing. Aruba recommends that baseline assessments of the data throughput of the organization be gathered to use in the mobility controller selection process. If a WLAN is already in place, use the AWMS before the Aruba WLAN is installed to help understand the average user throughput from wireless devices. If a WLAN is not currently in place, the network management system in place should be used to understand the size of traffic flows in the system.
- **Mobility controller role:** The role of the mobility controller in the system greatly affects the selection, because a master mobility controller has different requirements than a local mobility controller on the campus or a local mobility controller terminating RAPs or VIA clients. The most critical aspect to consider for master mobility controllers is the control processing power. However, local mobility controllers have greater concerns around data throughput and AP and user scaling.

Aruba recommends that information be gathered on a site-level basis during the planning process so that better choices are made for each site. Use the following table to keep track of this information.

Table 21 *Planning Guide*

Metric	Campus	Remote	VIA/VPN
AP Count			N/A
AM Count			N/A
RAP Count			N/A
Device Count			
VIA User Count	N/A	N/A	
VIA SSL Fallback?	N/A	N/A	
Avg. Data Throughput			

Controller Selection Formula – Local Controllers

Use the information gathered in the previous section to help determine the number and type of mobility controllers needed to meet the network goals of the organization. The local mobility controller tables that follow can be used for the majority of deployments, and will result in a correct mobility controller selection. A set of conditions is attached to each deployment model. If the deployment model fits within those conditions, the table should be used for mobility controller selection.

If the site has more devices or more APs than a single mobility controller can handle, increase the number of mobility controllers until sufficient capacity is attained. When redundancy is enabled, the number of controllers must be increased to account for redundant mobility controllers. The redundancy calculations are available at the end of this section.

Controller Scalability Table

The following table summarizes the key factors in selecting the proper mobility controller for the network.

Table 22 *Controller Scalability Table*

	Controllers							
Features	620 Controller	650 Controller	651 Controller	Aruba 3200	Aruba 3400	Aruba 3600	M3 Blade	Fully Loaded Chassis (4 x M3)
Max number of campus-connected APs per controller	8	16	17	32	64	128	512	2048
Max number of RAPs per controller	32	64	64	128	256	512	1024	4096
MAC addresses	2048	2048	2048	64000	64000	64000	64000	256000
Max number of users or devices per controller	256	512	512	2048	4096	8192	8192	32768
Maximum number of concurrent tunnels	256	512	512	2048	4096	4096	4096	16384

Table 22 *Controller Scalability Table (Continued)*

Features	Controllers							
	620 Controller	650 Controller	651 Controller	Aruba 3200	Aruba 3400	Aruba 3600	M3 Blade	Fully Loaded Chassis (4 x M3)
Max number of VIA clients per controller (no SSL fallback)	256	512	512	2048	4096	4096	4096	16384
Max number of VIA clients per controller (with SSL fallback)	128	256	256	1024	2048	2048	2048	8192
Maximum number of VLAN IP interfaces	128	128	128	128	256	512	1400	5600
Maximum firewall throughput	800 Mbps	2 Gbps	2 Gbps	3 Gbps	4 Gbps	4 Gbps	20 Gbps	80 Gbps
Maximum encrypted throughput (3DES, AESCBC256)	400 Mbps	1.6 Gbps	1.6 Gbps	1.6 Gbps	4 Gbps	8 Gbps	8 Gbps	32 Gbps
Maximum encrypted throughput (AES-CCM)	320 Mbps	800 Mbps	800 Mbps	800 Mbps	2 Gbps	4 Gbps	4 Gbps	16 Gbps

NOTE

This table and those that follow contain the maximum supported values for the mobility controllers. As with any other piece of networking equipment, caution should be exercised with any system that is approaching the maximum supported load. Aruba does not recommend that devices be run at full capacity except in extreme circumstances.

Local Mobility Controller – Campus or Branch Deployment

In a campus or branch deployment, the two most important factors are the required number of CAPs and AMs, and the required number of users on the campus. To select the proper controller, simply select the number of users on the site and the number of APs generated by VisualRF Plan or a traditional site survey.

Table 23 *Local Mobility Controller – Campus Deployment*

		CAP Count						
		8	16	17	32	64	128	512
Device Count	256	620 Controller	650 Controller	651 Controller	Aruba 3200	Aruba 3400	Aruba 3600	M3
	512	650 Controller	650 Controller	651 Controller	Aruba 3200	Aruba 3400	Aruba 3600	M3
	2048	Aruba 3200	Aruba 3200	Aruba 3200	Aruba 3200	Aruba 3400	Aruba 3600	M3
	4096	Aruba 3400	Aruba 3400	Aruba 3400	Aruba 3400	Aruba 3400	Aruba 3600	M3
	8192	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	M3

Local Mobility Controller - Remote Access Point Deployment

Selecting a mobility controller for a RAP deployment is more complex than selecting a campus mobility controller. The selection process is complex, because the RAP acts as a user of the system when it sets up the IPsec connection, and there are overall limits to the number of IPsec sessions on the system.

Table 24 *Local Mobility Controller – Remote Deployment*

		Recommended Controller for RAPs (Not Redundant)										
		RAP Count										
		1	2	4	8	16	32	64	128	256	512	1024
Total Number of Devices	1	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	650 Controller/ 651 Controller	Aruba 3200	Aruba 3400	Aruba 3600	Aruba 3600 x 2
	2	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	650 Controller/ 651 Controller	Aruba 3200	Aruba 3400	Aruba 3600	Aruba 3600 x 2
	4	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	650 Controller/ 651 Controller	Aruba 3200	Aruba 3400	Aruba 3600	Aruba 3600 x 2
	8	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	650 Controller/ 651 Controller	Aruba 3200	Aruba 3400	Aruba 3600	Aruba 3600 x 2
	16	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	650 Controller/ 651 Controller	Aruba 3200	Aruba 3400	Aruba 3600	Aruba 3600 x 2
	32	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	650 Controller/ 651 Controller	Aruba 3200	Aruba 3400	Aruba 3600	Aruba 3600 x 2
	64	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	650 Controller/ 651 Controller	Aruba 3200	Aruba 3400	Aruba 3600	Aruba 3600 x 2
	128	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	620 Controller	650 Controller/ 651 Controller	Aruba 3200	Aruba 3400	Aruba 3600	Aruba 3600 x 2
	256	650 Controller/ 651 Controller	650 Controller/ 651 Controller	650 Controller/ 651 Controller	650 Controller/ 651 Controller	650 Controller/ 651 Controller	650 Controller/ 651 Controller	650 Controller/ 651 Controller	Aruba 3200	Aruba 3400	Aruba 3600	Aruba 3600 x 2
	512	Aruba 3200	Aruba 3200	Aruba 3200	Aruba 3200	Aruba 3200	Aruba 3200	Aruba 3200	Aruba 3200	Aruba 3400	Aruba 3600	Aruba 3600 x 2
	1024	Aruba 3200	Aruba 3200	Aruba 3200	Aruba 3200	Aruba 3200	Aruba 3200	Aruba 3200	Aruba 3200	Aruba 3400	Aruba 3600	Aruba 3600 x 2
	2048	Aruba 3400	Aruba 3400	Aruba 3400	Aruba 3400	Aruba 3400	Aruba 3400	Aruba 3400	Aruba 3400	Aruba 3400	Aruba 3600	Aruba 3600 x 2
	4096	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600 x 2
	8192	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600	Aruba 3600 x 2

Aruba does not typically recommend the M3 mobility controller in remote access deployments. The reasons for this are practical (cost of the mobility controller and interface compatibility) as well as operational (increased user to RAP ratio as well as smaller failure domain with fewer RAPs on the Aruba 3600). M3s can be deployed as the RAP mobility controller. Simply replace the Aruba 3600s with M3s in the table. The quantity of each should remain the same.

Local Mobility Controller - VIA Deployments

For VIA user support, consider two factors: the number of devices and the use of SSL fallback. The following table shows the supported mobility controller based on device count, with and without SSL fallback enabled.

Table 25 *Local Mobility Controller – VIA Deployment*

Mode	620 Controller	650 Controller	651 Controller	Aruba 3200	Aruba 3400	Aruba 3600	M3 Blade
Max number of VIA clients per controller (no SSL fallback)	256	512	512	2048	4096	4096	4096
Max number of VIA clients per controller (with SSL fallback)	128	256	256	1024	2048	2048	2048

Calculating RAP and VIA Clients on the Same Mobility Controller

Aruba recommends that VIA and RAP deployments are separated onto different mobility controllers to simplify configuration, deployment, and troubleshooting. Determining the number of supported VIA clients depends greatly on the configuration of SSL fallback as well as the number of RAPs on mobility controller. Each RAP counts against three variables: the total RAP count, total user count, and the total IPsec tunnel limit. VIA clients count against the tunnel limit as well, but in instances where SSL fallback is enabled, two tunnels must be constructed for each VIA client. The formula for the mobility controller selection is:

> **Number of RAPs + Number of VIA Clients (x2 for SSL fallback) <= Mobility Controller IPsec Tunnel Limit**
>
> **Example:** A Aruba 3600 is being used as a remote access mobility controller and it has a maximum tunnel count of 4096 and a maximum of 8,192 users. With a full load of 512 RAPs, the mobility controller still has the capacity to terminate 3,584 VIA clients without SSL fallback, or 1,792 VIA clients with SSL fallback enabled. The total user count available on the RAPs depends on the number of VIA clients connected. The maximum is 3,584 (8,192 – 1,024 used by RAPs – 3,584 VIA clients) and 5,376 (8,192 – 1,024 used by RAPs – 1,792 VIA clients). That means an average of 7 and 10.5 users per RAP respectively.

Controller Selection Formula – Master Mobility Controller

The primary consideration for the master mobility controller is the scale of control-plane processing. The following table summarizes the capabilities of the mobility controllers when they act as the master in a mobility controller cluster without APs or users terminating directly on the master.

Table 26 *Master Mobility Controller Scalability*

Master	Maximum APs	Maximum Users or Devices
M3/Aruba 3600	4500	15000
Aruba 3400	2250	7500
Aruba 3200	1500	4500
650 Controller/651 Controller	250	1000
620 Controller	125	500

The M3 and Aruba 3600 have equivalent scalability numbers when they operate as the master mobility controller, so Aruba typically recommends that the Aruba 3600 be selected as the master for large-scale deployments.

Redundancy Considerations for Controller Count

The following table can be used to calculate the number of mobility controllers needed to provide a given level of redundancy.

Table 27 *Redundancy Planning*

Redundancy Model	Controller Count	Multiplier	Total
Master redundancy		X2	
Active-active		X2	
Active-standby		X2	
Many-to-one		Divide controller count by backup ratio (e.g. 3-1, divide by 3, 4-1 divide by 4, etc.)	
Full data center redundancy	Multiply all counts in the Total column by 2 to provide for full data center redundancy.		

When to Consider a Mobility Controller Upgrade

Over time, as the network becomes more utilized and the user population increases, the demands on the mobility controller will also increase. The following tables summarize the steps and commands that the network administrator should perform to judge the current capacity of the mobility controller. Some of these steps are available on the mobility controller, and others are accessed through the AWMS.

Mobility Controller

The following table describes the commands that can be issued from the CLI to examine the current state of the system. In most cases, the AWMS can be used to see state information over a longer time line.

Table 28 *Mobility Controller Monitoring*

<table>
<thead>
<tr><th>Command</th><th>Description</th></tr>
</thead>
<tbody>
<tr>
<td><b>Memory Utilization</b></td>
<td>
Memory is another limited resource on the system. When the system boots, it uses a set amount of memory to load ArubaOS and provide base-level functionality. Use the following command to show the current state of memory on the system:

```
(M3) # show memory

Memory (Kb): total: 1541620, used: 275960, free:
1265660
```

The organization should consider that the memory is moderately utilized at 30 Mb free (begin monitoring regularly) and highly utilized at 15 Mb free (being investigating) over a 5-minute period.

Alternatively, consider using AWMS to track average memory utilization over time.
</td>
</tr>
<tr>
<td><b>CPU Utilization</b></td>
<td>
As with all systems, a finite amount of CPU is available for processing data. Use the following command to find out the current CPU utilization of the mobility controller:

```
(M3) # show cpu

user 4.2%, system 2.8%, idle 93.0%
```

The organization should consider the CPU to be highly utilized at 70% (begin monitoring regularly) and considered critical at 100% (being investigating) over a 5-minute period.

Alternatively, consider using AWMS to track average CPU utilization over time.
</td>
</tr>
<tr>
<td><b>Platform Limitations for Users (Devices)</b></td>
<td>
As the platform reaches its maximum user count, the users will need to be split across multiple mobility controllers. During busy times of the day, use the following command to show the summary user count:

```
(M3) # show user-table | include Entries:

User Entries: 156/156
```

Compare the summary count to the platform limit. Or, consider using AWMS to track average user counts over time. The output would look like:

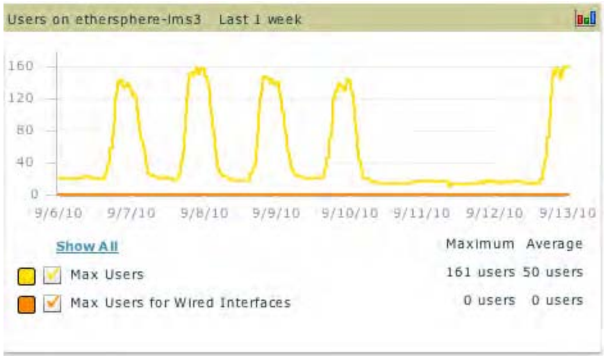

</td>
</tr>
</tbody>
</table>

Table 28 *Mobility Controller Monitoring (Continued)*

<table>
<tr><th>Command</th><th>Description</th></tr>
<tr><td><b>License Limitations</b></td><td>

License limitations can show up as the inability to add additional APs, or APs not behaving as expected due to incorrect license counts. Use the following command to ensure that all licensing numbers for APs match, and that sufficient license capacity exists on the platform:

```
(M3) # show license limits
```

If the platform is at its maximum capacity, additional mobility controllers must be purchased.

</td></tr>
<tr><td><b>Datapath Utilization / Throughput</b></td><td>

Mobility controllers have a limitation on the amount of traffic that can flow through the system. Use the following command to show this information as a percentage of utilization:

```
(M3) # show datapath utilization

Datapath Network Processor Utilization
------+---------+---------+----------+
      | Cpu utilization during past  |
  Cpu |  1 Sec     4 Secs    64 Secs |
------+---------+---------+----------+
    8 |      0% |      0% |       0% |
    9 |      0% |      0% |       0% |
   10 |      0% |      0% |       0% |
   11 |      0% |      0% |       0% |
   12 |      0% |      0% |       0% |
   13 |      0% |      0% |       0% |
   14 |      0% |      0% |       0% |
   15 |      0% |      0% |       0% |
   16 |      0% |      0% |       0% |
   17 |      0% |      0% |       0% |
   18 |      0% |      0% |       0% |
   19 |      0% |      0% |       0% |
   20 |      0% |      0% |       0% |
   21 |      0% |      0% |       0% |
   22 |      0% |      0% |       0% |
   23 |      0% |      0% |       0% |
   24 |      0% |      0% |       0% |
   25 |      0% |      0% |       0% |
   26 |      0% |      0% |       0% |
   27 |      0% |      0% |       0% |
   28 |      0% |      0% |       0% |
   29 |      0% |      0% |       0% |
   30 |      0% |      0% |       0% |
   31 |      0% |      0% |       0% |
```

The organization should consider that the data path is moderately utilized at 50% (begin monitoring regularly) and highly utilized at 70% (being investigating) over a 5-minute period.

Alternatively, consider using AWMS to track average data throughput over time.

</td></tr>
</table>

AWMS

On the AirWave Wireless Management Suite (AWMS) the information in the preceding table is grouped into a set of tables to show trending of data over time. This summary does not include license limits, but those can be viewed via the licenses link on the page. Use the numbers from the preceding table to calculate the remaining capacity on the mobility controller.

Figure 84 *AirWave Wireless Management Suite mobility controller summary*

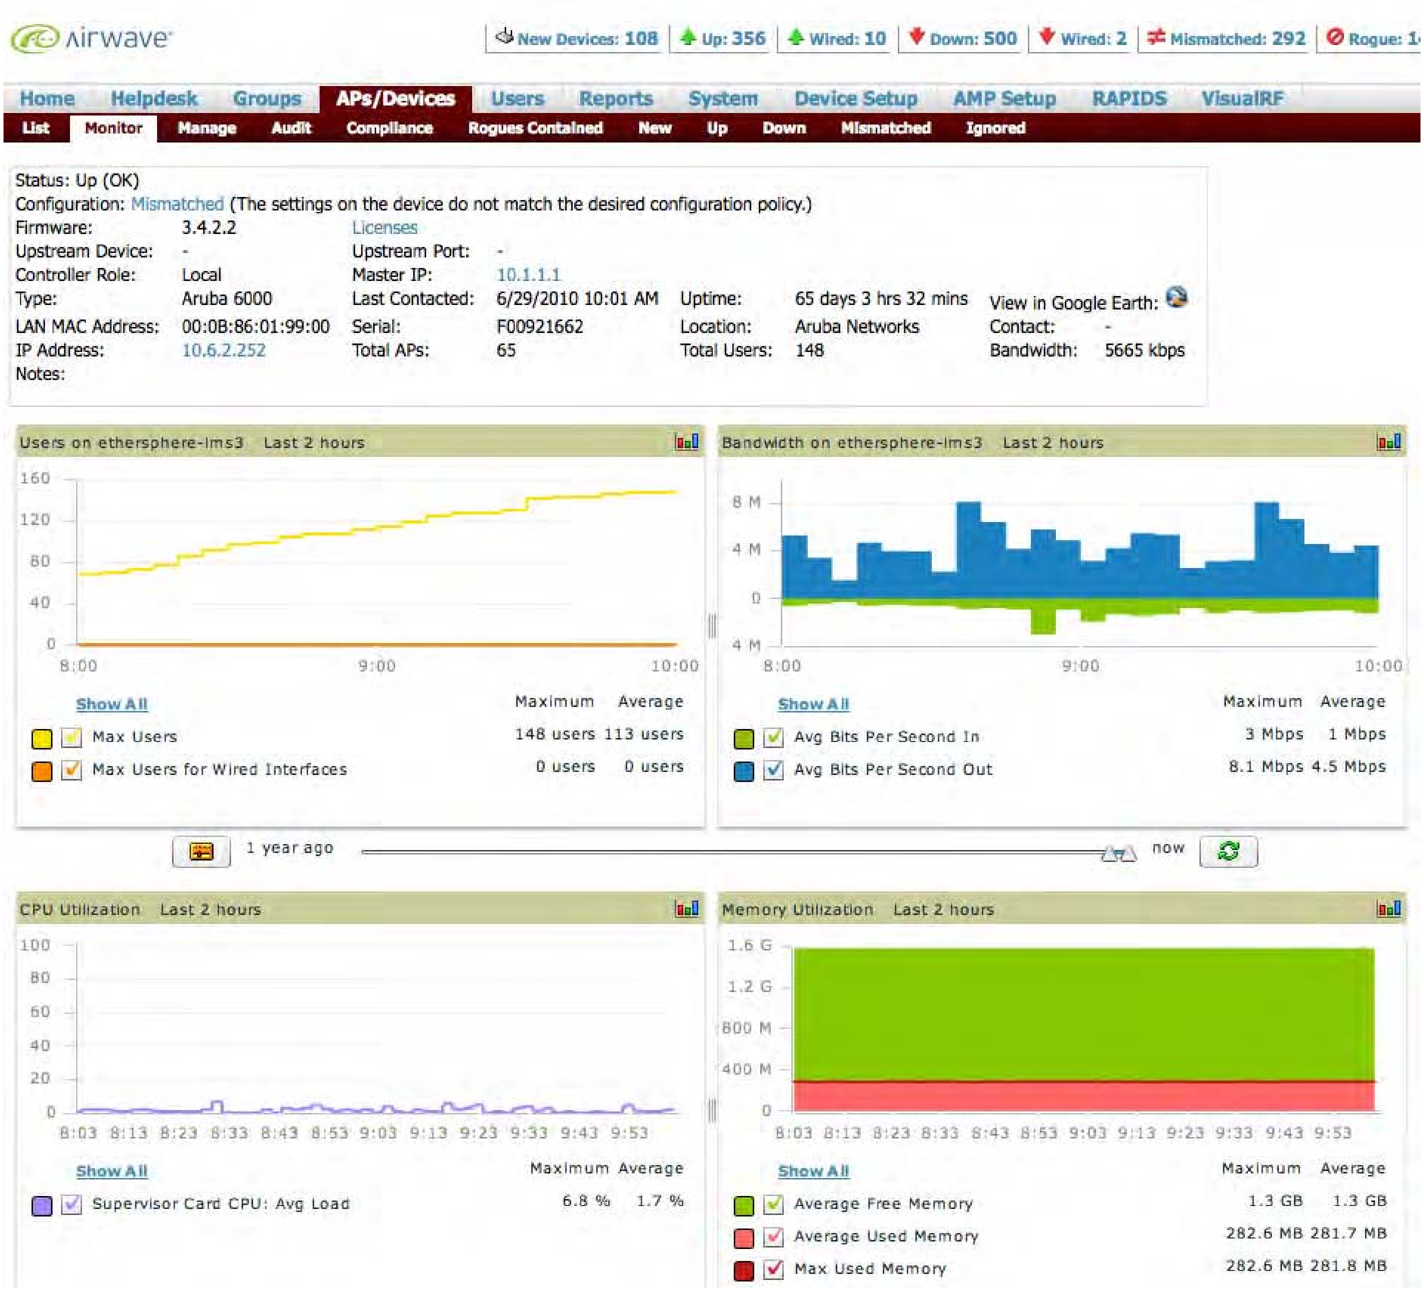

Status: Up (OK)
Configuration: Mismatched (The settings on the device do not match the desired configuration policy.)

Firmware:	3.4.2.2	Licenses					
Upstream Device:	-	Upstream Port:	-				
Controller Role:	Local	Master IP:	10.1.1.1				
Type:	Aruba 6000	Last Contacted:	6/29/2010 10:01 AM	Uptime:	65 days 3 hrs 32 mins	View in Google Earth:	
LAN MAC Address:	00:0B:86:01:99:00	Serial:	F00921662	Location:	Aruba Networks	Contact:	-
IP Address:	10.6.2.252	Total APs:	65	Total Users:	148	Bandwidth:	5665 kbps
Notes:							

	Maximum	Average
🟨 ☑ Max Users	148 users	113 users
🟧 ☑ Max Users for Wired Interfaces	0 users	0 users

	Maximum	Average
🟩 ☑ Avg Bits Per Second In	3 Mbps	1 Mbps
🟦 ☑ Avg Bits Per Second Out	8.1 Mbps	4.5 Mbps

	Maximum	Average
🟪 ☑ Supervisor Card CPU: Avg Load	6.8 %	1.7 %

	Maximum	Average
🟩 ☑ Average Free Memory	1.3 GB	1.3 GB
🟥 ☑ Average Used Memory	282.6 MB	281.7 MB
🟥 ☑ Max Used Memory	282.6 MB	281.8 MB

Use the graphs to track the following metrics from the mobility controller metrics table:

- **Memory:** Use the memory utilization graph to track free memory
- **CPU:** Use the CPU utilization graph to track CPU utilization
- **User count:** Use the Users on [device] graph to track users against maximum platform limits
- **Datapath utilization/throughput:** Use the Bandwidth on [device] graph to track datapath utilization in the system.

To track device utilization vs. license limits, click the Licenses link and note the license limits on the mobility controller. Next, scroll down the summary page and look for the summary of APs connected to the mobility controller. Compare this number vs. licensed limit and platform limits.

Figure 85 *AP count per mobility controller*

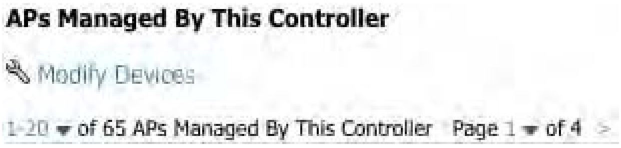

APs Managed By This Controller

Modify Devices

1-20 ▾ of 65 APs Managed By This Controller Page 1 ▾ of 4 >

An alternative method to track license exhaustion is to setup triggers in the system to alert when the system is reaching critical levels. Triggers are available for CPU, memory, and bandwidth. The following graphic shows the setup for the CPU trigger.

Figure 86 *AWMS trigger set to alert on a CPU load*

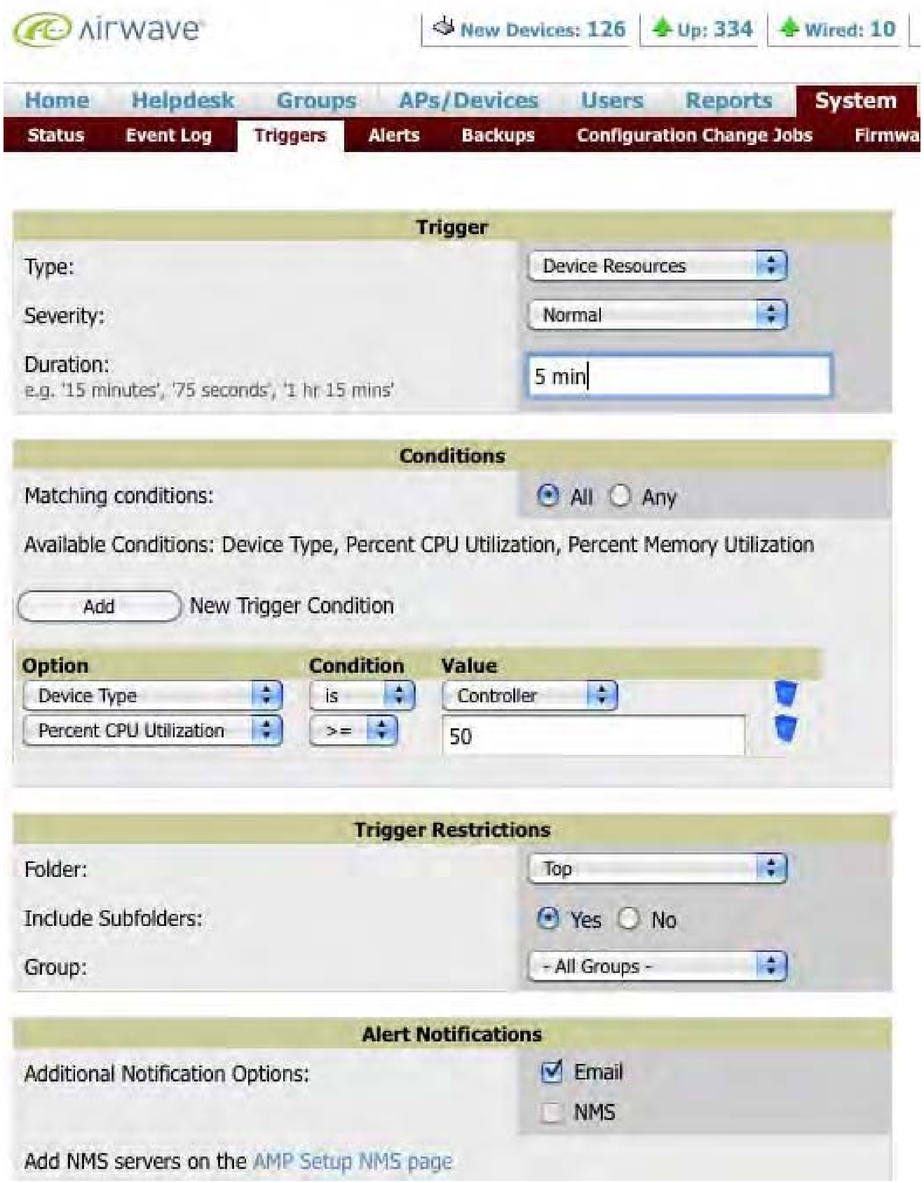

Trigger

Type:	Device Resources
Severity:	Normal
Duration: e.g. '15 minutes', '75 seconds', '1 hr 15 mins'	5 min

Conditions

Matching conditions: ● All ○ Any

Available Conditions: Device Type, Percent CPU Utilization, Percent Memory Utilization

Add New Trigger Condition

Option	Condition	Value
Device Type	is	Controller
Percent CPU Utilization	>=	50

Trigger Restrictions

Folder:	Top
Include Subfolders:	● Yes ○ No
Group:	- All Groups -

Alert Notifications

Additional Notification Options: ☑ Email ☐ NMS

Add NMS servers on the AMP Setup NMS page

Chapter 12

Putting It All Together

The components of an Aruba system can work independently to support single-purpose deployments, such as remote access or campus deployments. But the true power of the system comes through when the components are combined to provide seamless mobility and a consistent user experience no matter how the user connects. When Aruba products are deployed across the distributed enterprise network, the user can roam to any location and connect back without a change to their daily workflow. Users who can roam freely can be more productive and make fewer calls to the help desk.

Network administrators have only one infrastructure to learn and to maintain, which reduces the complexity and allows administrators to focus on other projects in the organization. The network team can do more with the same resources, which adds increased value to the organization.

Figure 87 *Distributed enterprise network*

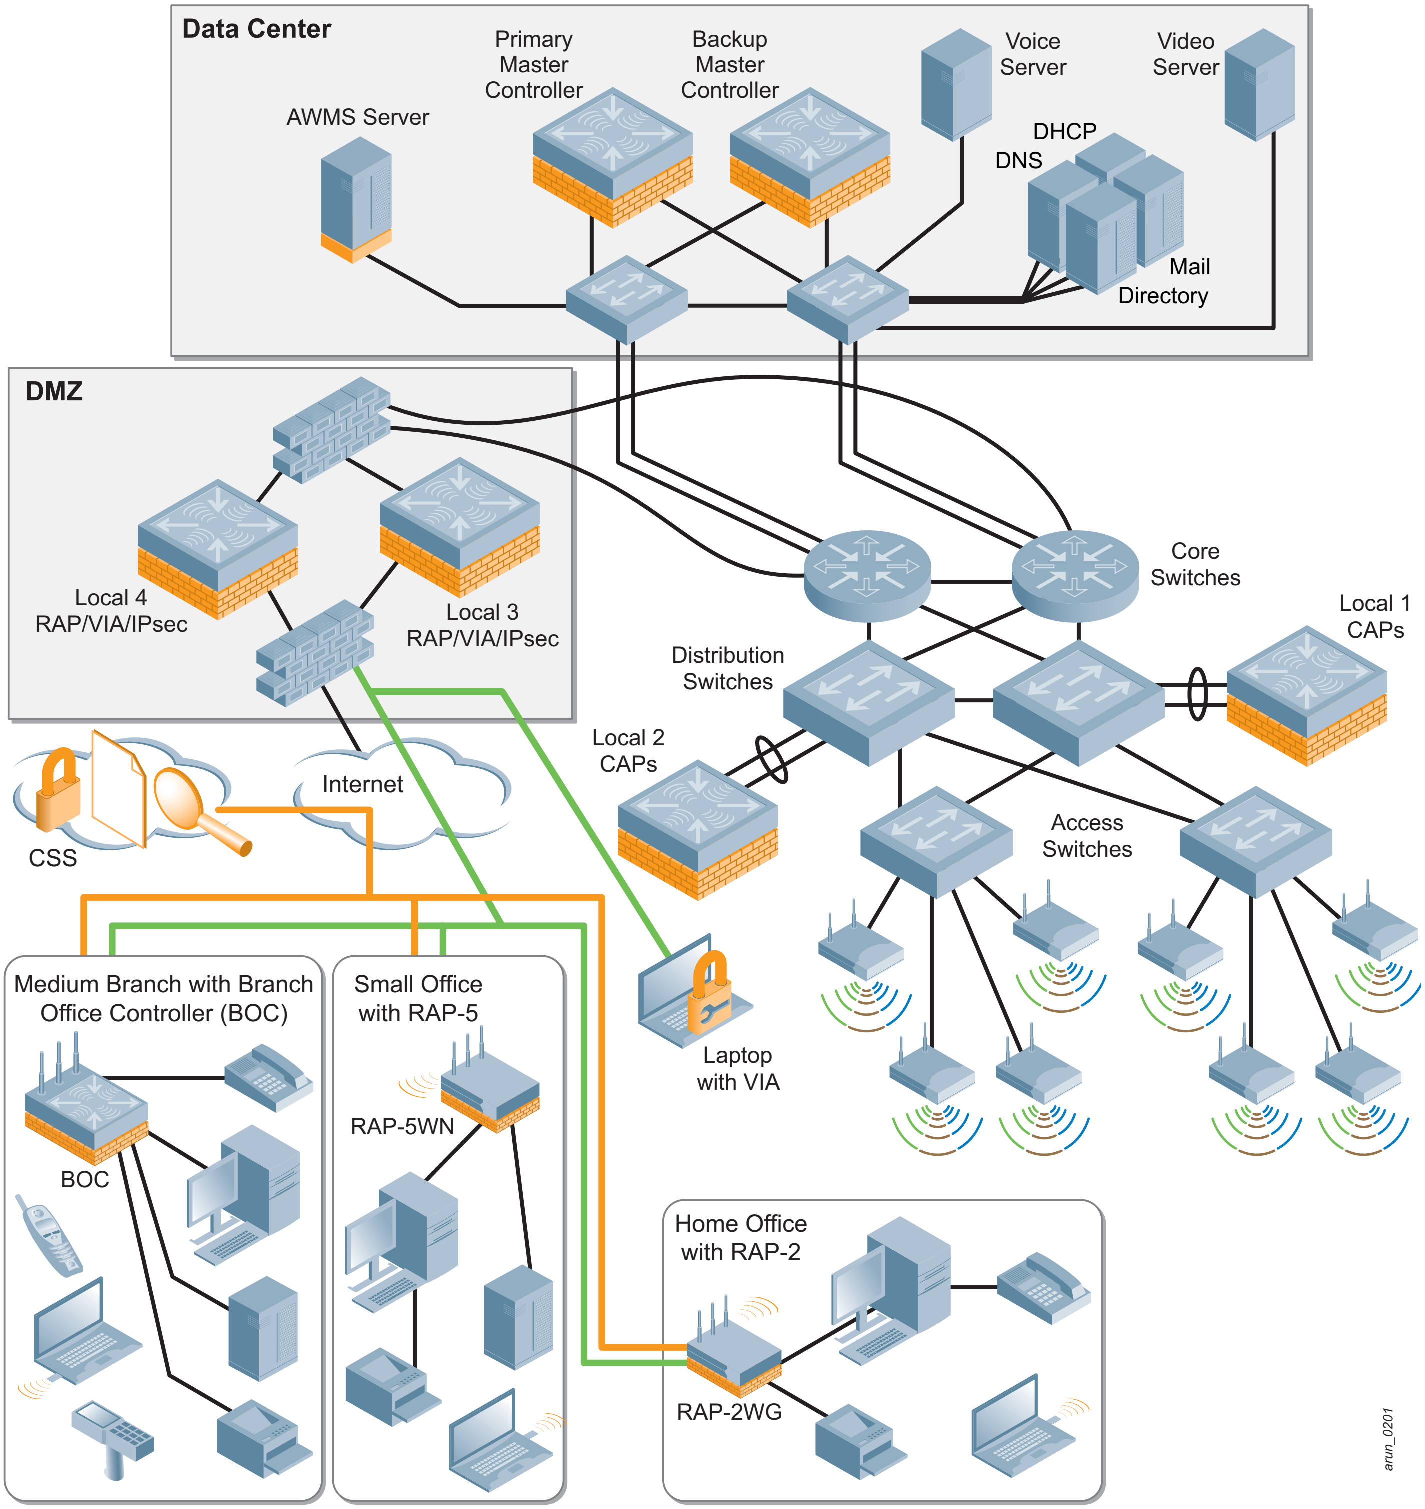

In this figure, all of the Aruba components work together to provide coverage to the campus, branch, home office, and remote user. By deploying Aruba products across the distributed enterprise network, the organization focuses on the user and provides a scalable, available network that increases user productivity through seamless connectivity.

Appendix A

CPsec Scalability

A number of factors can affect the deployment of CPsec and the scalability of the system. The following table provides information about the testing of the solution that Aruba has performed.

Feature Tested	Result	
Scalability of the AP white list synchronization	White list testing was performed with 140 local mobility controllers sharing a white list with a single M3 master. Synchronization consisted of a Aruba 5000 AP white list and synchronization across all controllers took approximately 10 minutes.	
Root CA scalability testing	Testing involved a single trust anchor with 140 locals directly attached to the single master. This level of scalability is applicable for the all masters deployment.	
Certification of 802.11n APs	M3 + 512 AP-12x	9 min 30 sec
	M3+ 256 AP-12x	6 min 00 sec
Certification of legacy 802.11a/b/g APs	M3 + 512 AP-70	29 min 30 sec
	M3+ 256 AP-70	19 min 00 sec
Boot time with CPsec OFF, 802.11n APs	M3 + 512 AP-12x	4 min 40 sec
	M3+ 256 AP-12x	3 min 30 sec
Boot time with CPsec ON, 802.11n APs	M3 + 512 AP-12x	9 min 20 sec
	M3+ 256 AP-12x	5 min 50 sec
Boot time with CPsec OFF, legacy 802.11a/b/g APs	M3 + 512 AP-70	5 min 10 sec
	M3+ 256 AP-70	3 min 50 sec
Boot time with CPsec ON, legacy 802.11a/b/g APs	M3 + 512 AP-70	9 min 45 sec
	M3+ 256 AP-70	6 min 30 sec

Appendix B

AP Failover Times

The following table shows failover times for APs failing over in an Active-Active deployment. This table was developed using two test cases:

- **Test case 1**: The test starts with APs distributed evenly between two mobility controllers, and ends when all APs have completed their transition from the disconnected mobility controller to the remaining mobility controller.
- **Test case 2:** This test takes the first test case, and includes both AP's and clients failing over between the two mobility controllers, and represents a 'worst case scenario'. Here the test starts with AP's and clients evenly distributed between two mobility controllers and end when the last client connection is re-established on the remaining mobility controller. Actual client experience will vary between a few seconds and the maximum value stated.

Client re-authentication rate is affected by a variety of factors outside of the WLAN infrastructure. In this scenario the clients dominate the resultant failover times. Client re-authentication can vary considerably base on the following:

- Client supplicant – problems with the client supplicant can include slow authentication and non-cached credentials
- Client driver – the client NIC card is slow to recognize that the network connection has been interrupted and is again available, and takes longer than expected to attempt to reconnect to the network
- Authentication type – different authentication types have different speeds for re-authentication, as an example 802.1X if more involved than an open network
- Insufficient AAA infrastructure – when large numbers of client attempt to reconnect to the network at the same time the AAA infrastructure, such as RADIUS and LDAP servers, can become overwhelmed
- Insufficient backend infrastructure – bottlenecks in the internal infrastructure can lead to longer response times and dropped packets, creating longer authentication times

NOTE

These numbers are based on Active-Active redundancy, with half of the APs and users active on each mobility controller. For Active-Standby or N+1 redundancy, expect that failover times and client authentication can take 25% to 100% longer for each set of numbers. This is due to the fact that a greater number of APs and clients will need to fail over to the backup controller. As an example, where 256 APs would need to fail over in the largest test case, 512 APs will need to fail over in an Active-Standby model.

Active-Active			
Test Case 1		**Test Case 2**	
64 CAP<==>64 CAP	20s	1K User+ 256 CAP<==>1K User + 256 CAP	2m:55s
128 CAP<==>128 CAP	31s	2K User+ 256 CAP<==>2K User + 256 CAP	4m:35s
256 CAP<==>256 CAP	52s	4K User+ 256 CAP<==>4K User + 256 CAP	9m:15

Appendix C

Scalability in the All Masters Deployment Model

The following scaling numbers apply to the all master deployments in ArubaOS 5.0/AWMS 7.0:

- An all master deployment, where APs are spread across multiple mobility controllers but are in the same physical location, must be under the control of a single AWMS instance. Currently these deployments scale to 2500 devices. WMS offload must be enabled on the mobility controllers to allow the AWMS to manage valid AP lists.

The following limitations exist in an all master deployment:

- WMS offload must be enabled, but the WMS data is not shared between AWMS instances. Currently deployments are limited to 2500 devices.
- Currently, configuration cannot be synchronized across multiple AWMS instances. If multiple AWMS servers are required, their configurations must be kept in sync manually.
- Depending on polling intervals, it can take some time for the AWMS to relearn that users and APs have moved to a new master. Assume at least one polling cycle for state to be reflected on the AWMS.
- If the status of a device changes on the controller, but changes again before AWMS polls, the controller and AWMS may contain different state information. This situation can occur with classification, but is more likely with user status. If the user roams more than once between polls, the AWMS will have only the most recent status and will not have the complete trail.
- When a failover occurs, and a client that was on the failed controller roams to the new controller, the client disappears from AWMS until polling finds the client again. This means that the client is not in the user page on AWMS and is not on the heat map.
- Locations that use multiple APs spread across multiple masters result in a wider margin of error than when a single master/local cluster is enabled.
- VisualRF heat maps can take multiple polling cycles to update after APs fail from one master to the backup master.