Campus Wireless Networks Validated Reference Design Version 3.3

ARUBA networks

Design Guide

Copyright

© 2008 Aruba Networks, Inc. All rights reserved.

Trademarks

AirWave®, Aruba Networks®, Bluescanner®, For Wireless That Works®, Mobile Edge Architecture®, People Move. Networks Must Follow., RFProtect®, The All Wireless Workplace Is Now Open For Business, and The Mobile Edge Company® are trademarks of Aruba Networks, Inc. All rights reserved.

Legal Notice

The use of Aruba Networks, Inc. switching platforms and software, by all individuals or corporations, to terminate other vendors' VPN client devices constitutes complete acceptance of liability by that individual or corporation for this action and indemnifies, in full, Aruba Networks, Inc. from any and all legal actions that might be taken against it with respect to infringement of copyright on behalf of those vendors.

Warranty

This hardware product is protected by the standard Aruba warranty of one year parts/labor. For more information, refer to the ARUBACARE SERVICE AND SUPPORT TERMS AND CONDITIONS.

Altering this device (such as painting it) voids the warranty.

www.arubanetworks.com

1322 Crossman Avenue
Sunnyvale, California 94089

Phone: 408.227.4500
Fax 408.227.4550

Contents

Chapter 1

Introduction

This design guide is one of a series of books that describes Aruba's User-Centric Network Architecture and provides network administrators with guidelines to design and deploy a centralized enterprise-wide wireless LAN (WLAN) network for the most common customer scenarios.

This guide complements the technical documentation you received with software and hardware releases for Aruba components.

Aruba Reference Architectures

An Aruba Validated Reference Design (VRD) is a package of network decisions, deployment best practices, and detailed descriptions of product functionality that comprise a reference model for common customer deployment scenarios. The VRD presented in this guide is representative of a best practice architecture for a large Campus WLAN serving thousands of users spread across many different buildings joined by SONET, MPLS, or other high-speed, high-availability network backbone.

The Campus Wireless Network is one of five reference architectures commonly deployed by our customers. For a brief description of the other deployment models refer to Appendix C, "Alternative Deployment Architectures" on page 71.

Reference Documents

Refer to the following documentation for more detailed technical information about Aruba OS.

Title	Version
ArubaOS User Guide	3.3.1
ArubaOS CLI Guide	3.3.1
ArubaOS Release Note	3.3.1
ArubaOS Quick Start Guide	3.3.1
MMS User Guide	2.5
MMS Release Notes	2.5

Contacting Aruba Networks

Web Site Support	
Main Site	http://www.arubanetworks.com
Support Site	http://www.arubanetworks.com/support
Software Licensing Site	https://licensing.arubanetworks.com
Wireless Security Incident Response Team (WSIRT)	http://www.arubanetworks.com/support/wsirt
Support Email	support@arubanetworks.com
WSIRT Email Please email details of any security problem found in an Aruba product.	wsirt@arubanetworks.com

Telephone Support	
Aruba Corporate	+1 (408) 227-4500
FAX	+1 (408) 227-4550
Support United States France United Kingdom Germany All Other Countries	 800-WI-FI-LAN (800-943-4526) +33 (0) 1 70 72 55 59 +44 (0) 20 7127 5989 +49 (0) 69 38 09 77 22 8 +1 (408) 754-1200

Chapter 2

Aruba's User-Centric Network Architecture

This chapter provides an overview of a centralized wireless LAN architecture, followed by a high level technical overview of the Aruba User-Centric Network components and network design.

This overview describes the technology, architecture, services, and applications that make up an Aruba User-Centric Network to help you make the right design choices, and select the appropriate solution components.

Understanding Centralized Wireless LAN Networks

In the early days of wireless LAN (WLAN) networks, Access Points operated in an autonomous fashion much like other routers and switches in the network. Access Points were managed and maintained independently; which worked for very small wireless deployments, such as lobbies and conference rooms where guests were expected.

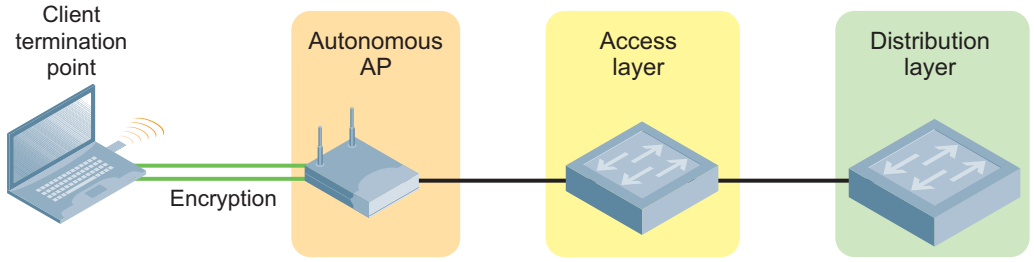

As large numbers of regular enterprise users began to expect connectivity using wireless connections, the autonomous Access Points became a management, reliability and security headache. Maintaining consistent configurations for dozens or hundreds of standalone APs became time-consuming, and introduced errors. Because each AP was a standalone device, network availability could not be guaranteed if any single AP failed. Centralized management consoles also fell short of expectations; and, in general, never grew beyond a certain point due to escalating operational costs. The workload associated with maintaining security, managing and troubleshooting large numbers of APs created a barrier to adoption in the larger enterprise; except in niche applications, such as guest access in conference rooms.

From a security perspective, users did not experience true mobility because network managers addressed WLAN security issues by treating wireless users and remote dial-up users the same way. Oftentimes, wireless users are quarantined on a single VLAN and forced through the "de-militarized zone" (DMZ) residing outside the corporate intranet. Users are then expected to tunnel into the corporate network through VPN concentrators that support industrial strength encryption such as AES.

A VPN was required primarily because of the 'port-based security' limitation of modern enterprise network infrastructures. VLANs and access controls are specified at the port level. When an autonomous AP is plugged in, then all users who connect to that AP inherit those security settings whether they are supposed to have them or not.

VPNs were a rudimentary way to impose identity-based authentication and provide extra encryption for first-generation wireless security systems. Unfortunately, these VPN concentrators were optimized for low speed WAN connections not intended for large numbers of high-speed wireless LAN users which then resulted in poor performance, management complexity, mobility, and scalability problems.

Introducing Aruba's User-Centric Network

In recent years, controller-based wireless switch architectures have been widely adopted to overcome the limitations of the autonomous AP. The Aruba centralized WLAN model shown below represents a structured model for WLAN deployment and ongoing management using a holistic approach to build enterprise WLANs that support user mobility without sacrificing security, manageability and scalability.

The Aruba User-Centric Network is an "overlay" network consisting of a centralized Mobility Controller and thin APs that work together over an existing high-speed network. Most enterprise networks have been engineered for high performance and high reliability, therefore, deploying the Aruba User-Centric Network as an overlay will not adversely affect the investment and reliability of the existing network.

With this approach, a centralized appliance controls hundreds or thousands of network-attached radios in a secure, reliable manner. This model represents a unified mobility solution integrating user mobility, identity based security, remote access, and enterprise fixed mobile convergence (eFMC) solutions.

Centralized WLAN Model

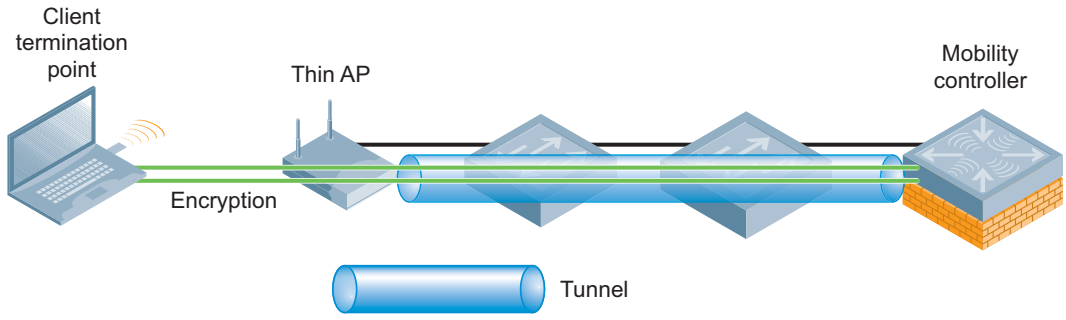

In this system, the intelligence that once resided in autonomous APs is now integrated into a centralized WLAN Mobility Controller designed for high-performance 802.11 packet processing, mobility and security management. These controllers are typically deployed in secured data center environment or distribution closets with redundant power and connectivity. APs are simplified and become network-attached radios that perform only transceiver and air monitoring functions. These access points are commonly referred to as "thin" APs. Connected to the Mobility Controller directly or over a layer 2/3 network by encrypted tunnels, they become extended access ports on the Mobility Controller directing user traffic to the controller for processing; while providing visibility and control of the RF environment to protect against intrusions (such as unauthorized users or rogue APs).

ArubaOS and Mobility Controller

This section describes Aruba's operating system features, optional add-on modules and the Mobility Controller that comprise Aruba's User-Centric Network Architecture.

ArubaOS

The ArubaOS serves as the operating system and application engine for all Aruba Mobility Controllers, and is the core component that enables user-centric networks. Standard with every Aruba Mobility Controller, ArubaOS provides unprecedented control over the entire mobile environment enabling Aruba's unique adaptive wireless LANs, identity-based security, and application continuity services.

The main features of ArubaOS include:

- Sophisticated authentication and encryption
- Protection against rogue wireless APs
- Seamless mobility with fast roaming
- Adaptive RF management and analysis tools
- Centralized configuration
- Location tracking and more

ArubaOS also offers the following optional add-on modules that provide advanced capabilities including wireless intrusion protection (WIP), identity-based security with user-centric policy enforcement, mobile Network Access Control (NAC), secure remote access, and advanced network connectivity technologies.

- Wireless Intrusion Protection
- Policy Enforcement Firewall
- VPN Server, Remote AP
- External Services Interface
- Voice Services Module
- Wireless Mesh, and xSec Advanced L2 Encryption.

A complete description of all software modules is available in Appendix A, "Licenses" on page 67 of this document.

Mobility Controller

The Aruba Mobility Controller is the center of the User-Centric Network. The Mobility Controller is a part of a purpose built, scalable appliance family that runs the ArubaOS operating system and software modules. It provides network administrators the ability to manage the system state and rapidly scope problems for individual users across a single Master/Local controller cluster in a network. Refer to the Aruba Mobility Management System (MMS) in Chapter 9, "Controller Clusters and the Mobility Management System™" on page 63 to manage more than one Master/Local Controller cluster.

The Mobility Controller provides advanced RF features that take guess work and maintenance out of maintaining a wireless LAN. With RF Plan, a predictive site survey can be performed with nothing more than a floor plan and coverage requirements. Once installed, the system's Adaptive Radio Management (ARM) takes over. This distributed and patented algorithm runs to constantly monitor the RF environment, and adjust AP power and channel settings without user intervention; even in the face of interference or AP failure. RF Live shows the actual real time coverage using "heat maps" overlaid on the floor plan, while RF Locate allows Wi-Fi® clients and active RFID tags to be triangulated on the same set of floor plans.

Once the RF is running, security is initiated. Aruba Mobility Controllers use a multi-layered system to provide continuous protection of the network. The system constantly scans the environment looking for threats to users, and takes proactive action to contain rogue access points and potential attackers. Strong encryption and authentication techniques are routinely used to ensure users can safely connect to the network and that all transmissions are secure. The Mobility Controller uses a stateful firewall to monitor client traffic for policy violations and to provide high touch services.

Now that RF is present and secure, users are ready to roam the enterprise. Aruba's IP Mobility feature provides the capability for users to roam the enterprise without losing their connection or changing their IP address, even when moving between APs or controllers. This is critical when the organization moves to Voice over WLAN and dual mode phones.

Multi-function Thin Access Points

Aruba's access points serve multiple functions depending on their role in the network. APs are either indoor or outdoor deployable; and are available with various options, such as fixed or removable antennas, single or dual radio APs, and depending on the AP, can operate in one or more of the a/b/g/n spectrums. Selection of hardware based options should be considered depending on the deployment.

Functionality is defined by the role assigned through software modules and administrator configuration. Each radio on an Aruba AP can serve in one of five different roles. These roles include:

- Access Point (Local AP)
- Air Monitor
- Mesh Portal
- Mesh Point
- Remote AP

In some modes, the Aruba APs can operate as remote capture devices saving the network staff from having to walk to a problem area to use a handheld sniffer for troubleshooting.

Access Point

The most typical deployment uses an Aruba AP in the Access Point role. In this role, the AP radio(s) are used to connect user to the network infrastructure. The AP acts as a thin radio with much of the functionality of the system taking place on the Mobility Controller. Traffic is not processed on the AP. Instead, it is tunneled as an encrypted 802.11 frame to the controller via GRE. When an AP is connected to access layer switches it is known as a "campus-connected" or "local" AP.

Air Monitor

Used as an Air Monitor, the AP works as a network sniffer. The air monitor looks for rogue APs, monitors the RF environment and wired environment, and when combined with the wireless intrusion detection system (WIDS) software license it acts as a WIDS sensor to protect the network from those violating policy. The system can classify interfering and rogue APs based on network traffic and RF monitoring. Aruba APs can be dedicated to the Air Monitor function or can perform this role on a part-time basis when configured in the Access Point role.

Aruba recommends using dedicated Air Monitors for deployments of latency sensitive applications such as voice and video. Typically, one Air Monitor can provide security to the area served by up to four Access Points.

Mesh Portal or Mesh Point

In the Mesh Portal or Mesh Point role, the AP is taking part in Aruba's secure enterprise mesh network. This network is based around a single AP (the Mesh Portal) with a wired network connection, and one or more Mesh Point APs performing wireless backhaul or bridging of network traffic.

When used with dual radio APs, the mesh devices can provide client access on one radio and backhaul on the second. User traffic is authenticated and protected by the same centralized encryption method as wired APs, while Control traffic is protected by WPA2 authentication and encryption.

Aruba's Secure Enterprise Mesh Network

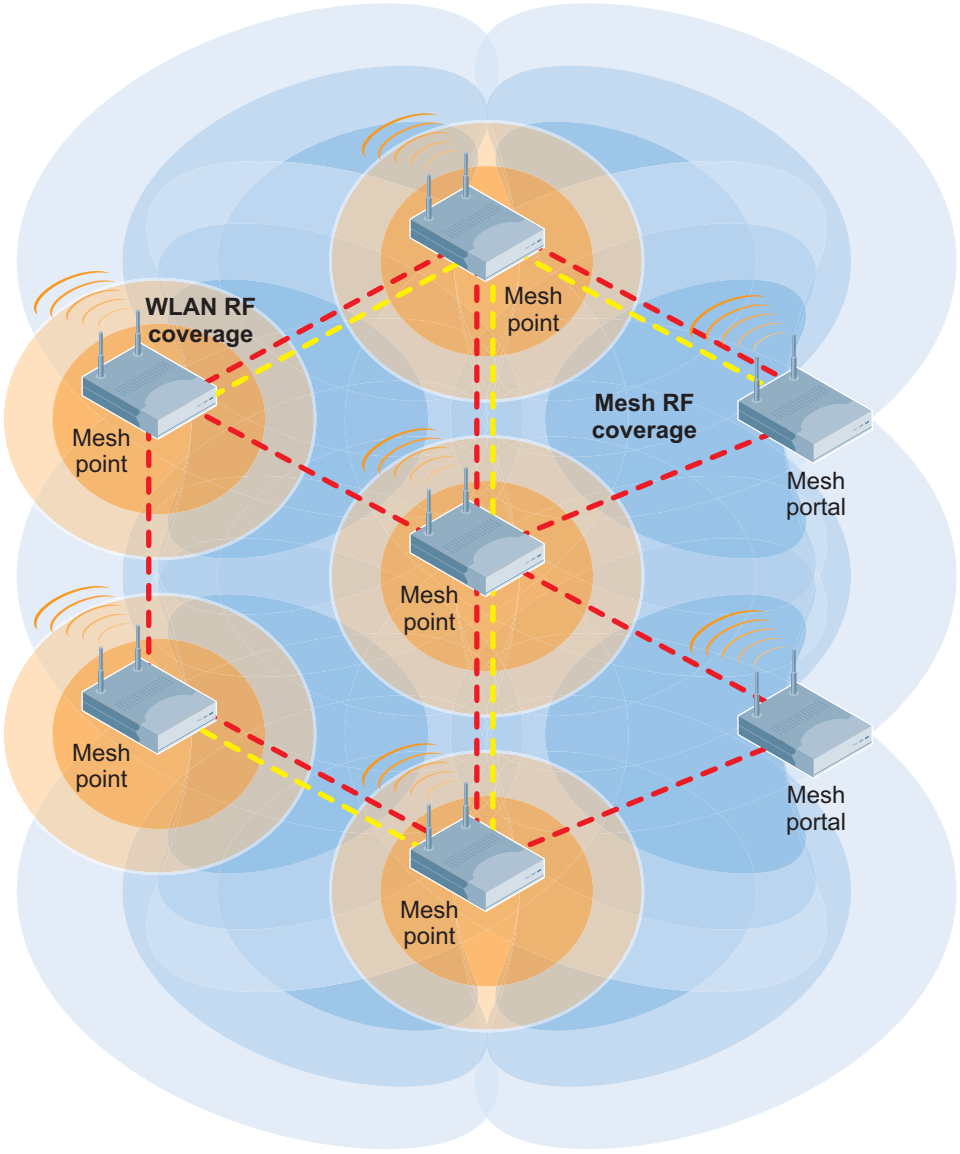

Remote AP

Using the Remote AP license, the AP can be used as a remote access device across a WAN. Plugging in to any Internet capable Ethernet port, the AP will create a secure tunnel using IPSec (AES) to a designated Mobility Controller. Typically this is done at corporate headquarters, or in regional data centers around the world for global deployments. The same SSIDs, authentication, and security are then available anywhere in the world.

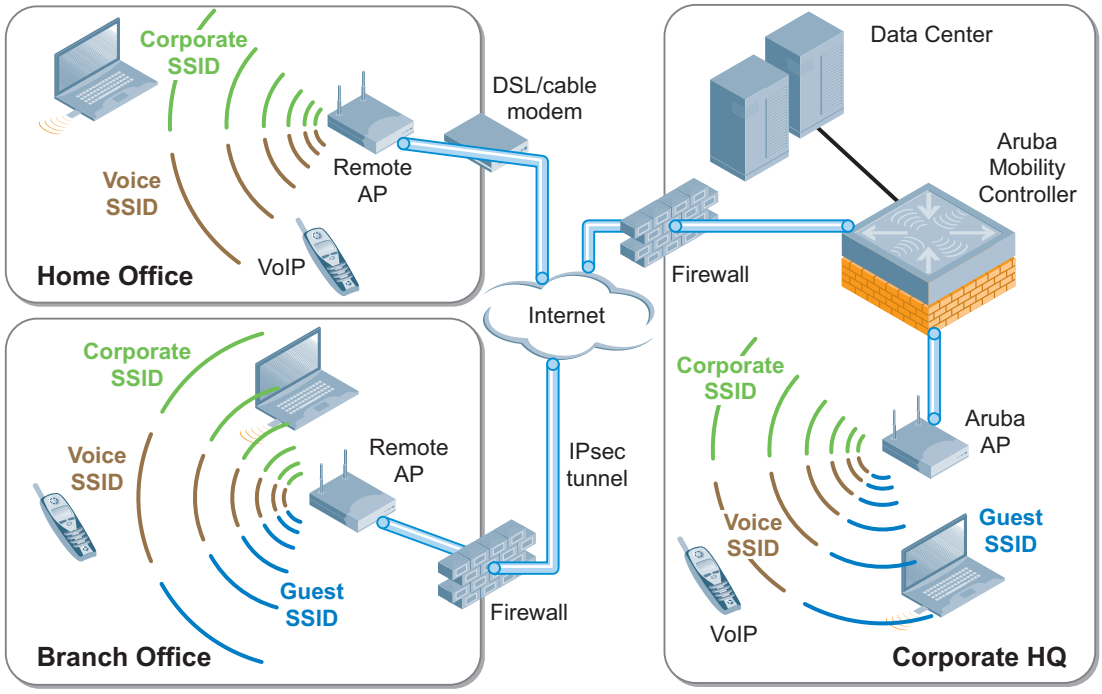

This provides an on-demand corporate hotspot with the same security and access to resources that users will find at the corporate campus without having to install additional software or be subject to a software learning curve. Unlike a software VPN that provides only a limited set of services, using the Aruba Remote AP license extends the entire corporate WLAN experience with the same powerful User-Centric Security.

Mobility Management System

Wireless networking doesn't make the IT administrator's job easier; in fact, it can make the job considerably harder. There are no longer any wires to trace, and IP address information only tells you where that user started their day. The MMS consists of a new set of tools to help administrators understand and visualize the wireless network they are administering. It is designed to provide network administrators with the ability to effectively manage multiple Master/Local clusters in the network. The user-centric management model allows administrators to rapidly visualize all network objects related to the user in real-time; drastically reducing the mean-to-resolution (MTTR) while ensuring a high quality WLAN user experience.

The Mobility Management System™ consists of a built-in location API that enables external systems to query the location of any WLAN device. The Mobility Management System software can be deployed on any PC platform (Linux or Windows 2003) or as an option, can be purchased as an enterprise class, hardened appliance.

One controller in each Aruba deployment is designated as the Master Controller. The Master Controller can also manage "Local" controller pairs, or clusters, in a high-availability configuration. However, once

the network grows to multiple clusters, a single centralized view across multiple Master/Local controllers of the following key operational data becomes highly desirable.

- Users on wireless network
- APs that users are connected to
- 802.11 traffic statistics
- AP failure notifications
- Failover alternatives and backup coverage maps

Mobility Management System

Refer to Chapter 9, “Controller Clusters and the Mobility Management System™” on page 63 for more detailed information.

To help set the stage for the complex campus network presented in Chapter 4 on page 19, it is useful to begin with a very small network. In this chapter, we consider a network that is typically deployed in a Proof-of-Concept (PoC) test involving a handful of Access Points and a Master Controller that provides guest and employee coverage to a conference room.

PoC Network - Physical Design

To keep the example as simple as possible, the design of this network involves a single AP and a single Mobility Controller, and uses an existing RADIUS or LDAP server for authentication.

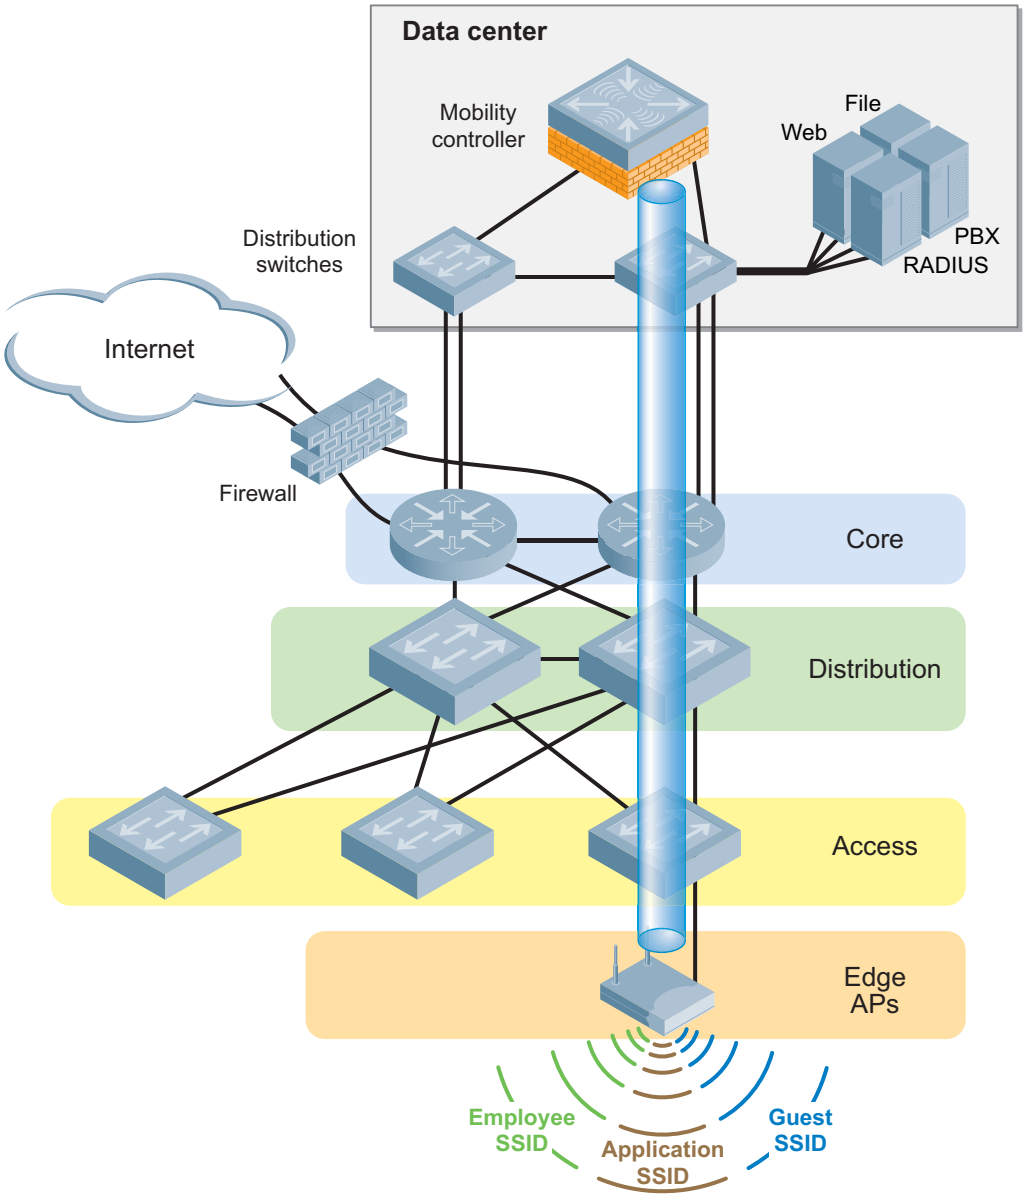

In this network, the AP has been deployed into a conference room, and is connected to the existing VLAN provided for wired users. In keeping with the concept of a network overlay, no reconfiguration or special VLANs need to be created as long as the access point has IP connectivity to the Mobility Controller.

PoC Network - Logical and RF Design

A common feature of centralized WLAN architectures is the ability to support many Service Set Identifiers (SSIDs) simultaneously from the same APs. Each SSID can have its own authentication and encryption settings based on the capabilities of the clients and the services that each needs. In this PoC network there are three SSIDs available for association via the AP.

- Employee
- Guest
- Application

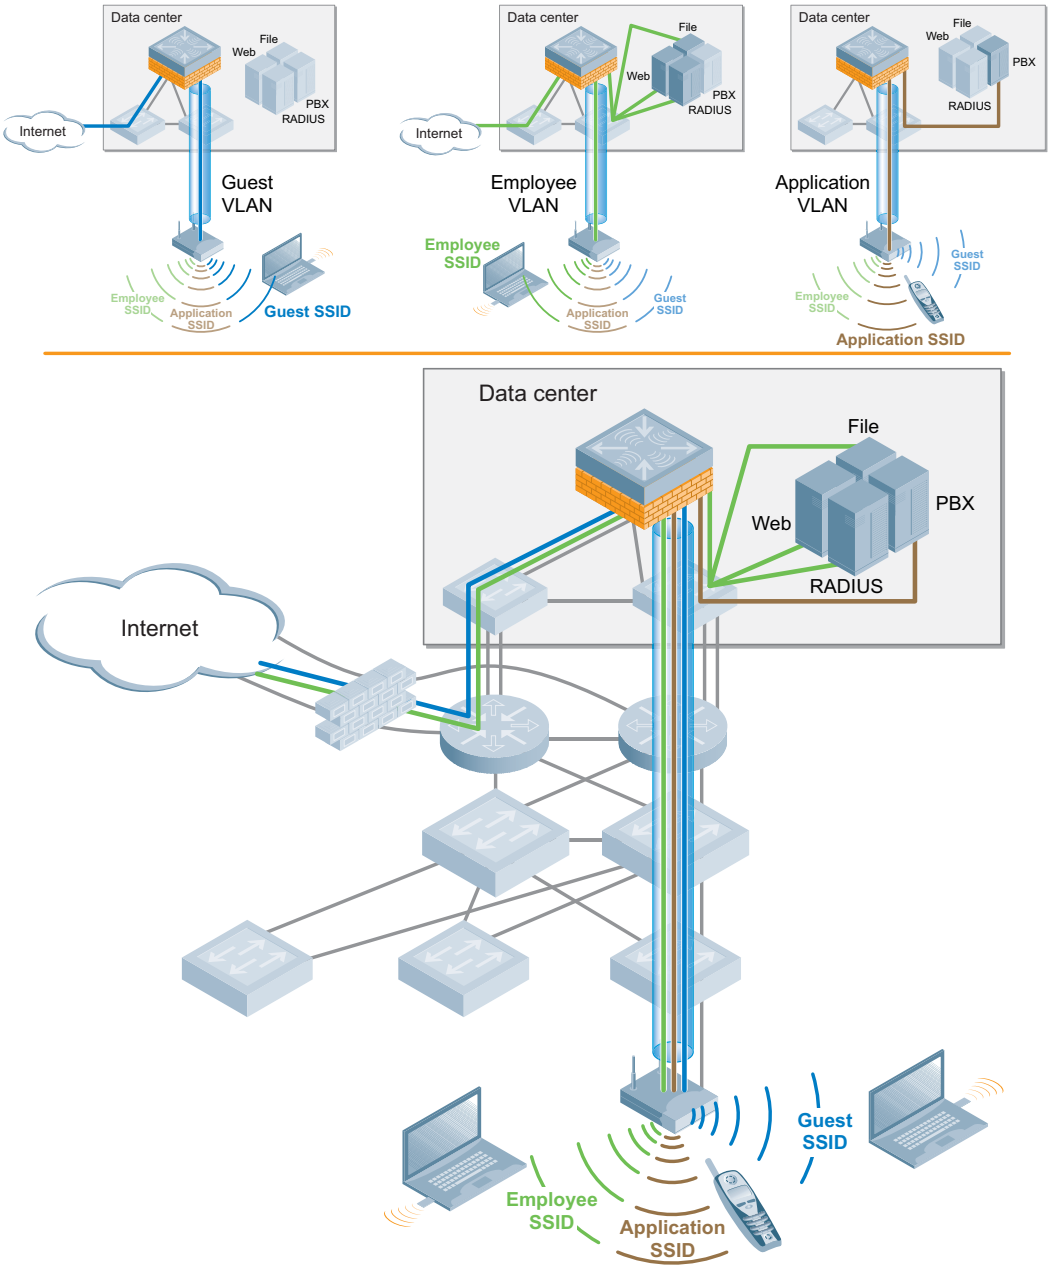

Users will associate to the Access Point and authenticate with the RADIUS server that already exists in the network. Employee users will use the Employee SSID, while guests will use the Guest SSID. Voice and data devices will associate to the Application SSID, and will be given a role based on the network services they are capable of accessing.

Each user and device has a specific role and associated policy enforced by the stateful firewall in the Mobility Controller. The Employee user now has full access to all resources within the network and the internet. Guest users are only permitted to access the Internet using specific protocols such as HTTP and HTTPS. Application devices are only able to access related application servers; for example, a phone running SIP can only access the SIP server to make calls.

This simple network describes the overlay functionality of an Aruba network, and shows how network control and policy enforcement is built into the fabric of the system. Users are only able to access those resources they have permissions for, and only after they have successfully authenticated to the network. This is the definition of an Aruba User-Centric Network.

Chapter 4

Campus WLAN Validated Reference Design

This chapter presents a more complex network model representing a common Aruba deployment in a large campus WLAN environment.

Enterprise networks support thousands of employees, with rigorous service level expectations. To meet these requirements, a reference wired network architecture that defines Core, Distribution and Access elements has become well established among IT network professionals. These elements form the building blocks of large scale, highly-available networks. Vendor validation of their products against this conceptual reference architecture provides IT organizations with assurance that products will perform and interoperate as expected.

Aruba User-Centric Enterprise Wireless Networks also support large numbers of users with stringent service level expectations. To enable IT network architects to successfully plan deployments, Aruba has developed a Validated Reference Design (VRD) that leverages the experience of more than 3,500 customer deployments, peer-review by Aruba engineers, and extensive performance testing. This reference design leverages and extends the familiar wired model in order to deploy a user-centric network as an overlay.

Aruba Campus WLAN Physical Architecture

The Validated Reference Design network model described in this chapter is referenced throughout the remainder of this book. The model depicts a cluster-based architecture typical of large enterprise deployments. For this type of deployment it is a best practice to employ distributed control and data planes using a hierarchical 'Master/Local' strategy with separate controller clusters providing each service. This will provide a scalable highly available architecture for data and voice traffic throughout the enterprise.

Some key components of this reference model include:

- Master Controllers – Two MMC-3600 model appliances configured to use Master redundancy. Each controller has redundant gigabit Ethernet links into the data center distribution switches, and share a Virtual Router Redundancy Protocol (VRRP) address.
- Local Controllers – Aruba Local Controllers consist of Multiservice Mobility Module blades in an MMC-6000 chassis. In the Aruba VRD, these Mobility Controllers are running in "active-active" redundancy, with two VRRP addresses shared between them. Each controller has two 10 gigabit Ethernet links bonded via Etherchannel to a single distribution layer switch.
- Access Points – Dual radio (A/B/G) AP65 access points are deployed throughout the enterprise carpeted space, providing high bandwidth access across the 2.4 GHz and 5Ghz bands. These APs are densely deployed. "Dense Deployment" uses a microcell architecture to cover an area using overlapping APs at relatively low transmit power. This design strategy enables ARM to detect and close coverage holes in the event of an AP failure by increasing power on neighboring APs. Smaller cells also help ensure proper load balancing of Voice over WLAN callers.
- SSIDs – There are three Service Set Identifiers present in the Reference Design. One SSID is used for employees and runs WPA2 for authentication and encryption. A second SSID is used by applications such as voice or video, and runs WPA with a Pre-Shared Key for authentication and encryption. The final SSID is open with a web based captive portal for authentication and is used by guests. Each user or device that associates with the network is placed in a role that is enforced by the stateful firewall.

- Air Monitors – AMs are deployed at a ratio of one AM for every four APs deployed. These handle many of the IDS related duties for the network, and will assist in drawing accurate heat maps displaying graphical RF data. Aruba considers dedicated Air Monitors to be a security best practice because they provide full time surveillance of the air.

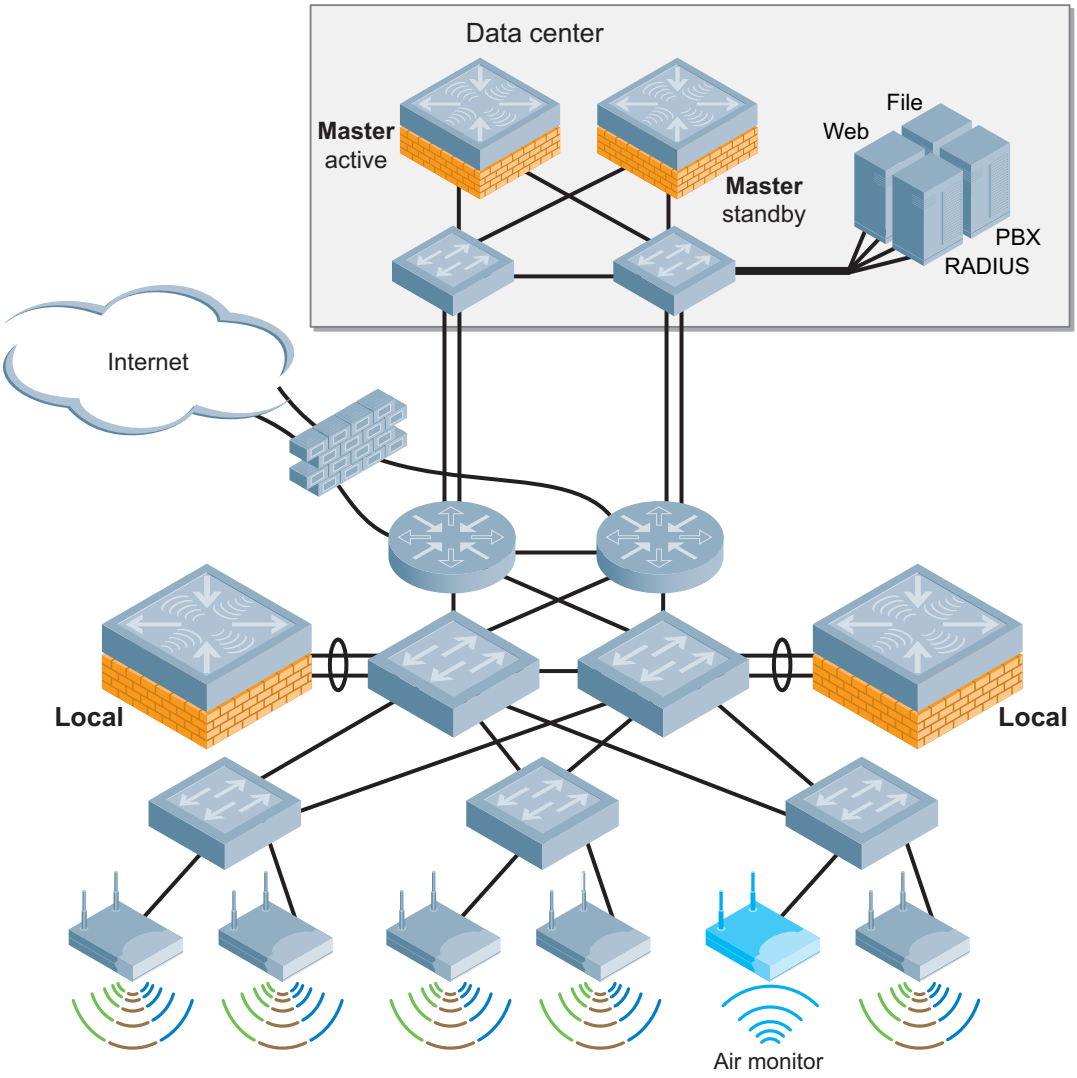

Aruba Campus WLAN Logical Architecture

From a logical perspective, the VRD overlay introduces three new terms into the familiar “core/ distribution/access” framework. They are “Management,” “Aggregation” and "Wireless Access."

- Management

  The Management layer provides a distributed control plane for the Aruba User-Centric Network that spans the physical geography of the wired network. Critical functions provided by the Management Layer Mobility Controllers include L3 client mobility across Aggregation layer controllers, and failover redundancy. Typically, larger networks, such as campus systems also off load ARM and IDS processing from the Aggregation Layer to the Management Layer.

- Aggregation

  The Aggregation layer is the interconnect point where wireless traffic is aggregated and enters or exits the wired network. Secure encrypted GRE tunnels from APs at the Wireless Access layer terminate on controllers at the Aggregation layer. This provides a logical point for enforcement of roles and policies, and is where the ArubaOS creates the User-Centric Network Experience.

Aggregation Layer Mobility Controllers allow user traffic to stay close to associated servers; there is no need to tunnel user traffic all the way to the Management layer.

- Wireless Access

  The Wireless Access layer is comprised of APs: single or dual-band, 802.11a/b/g or n, indoor or outdoor. They can be connected using wired switch ports, secure mesh or Remote AP.

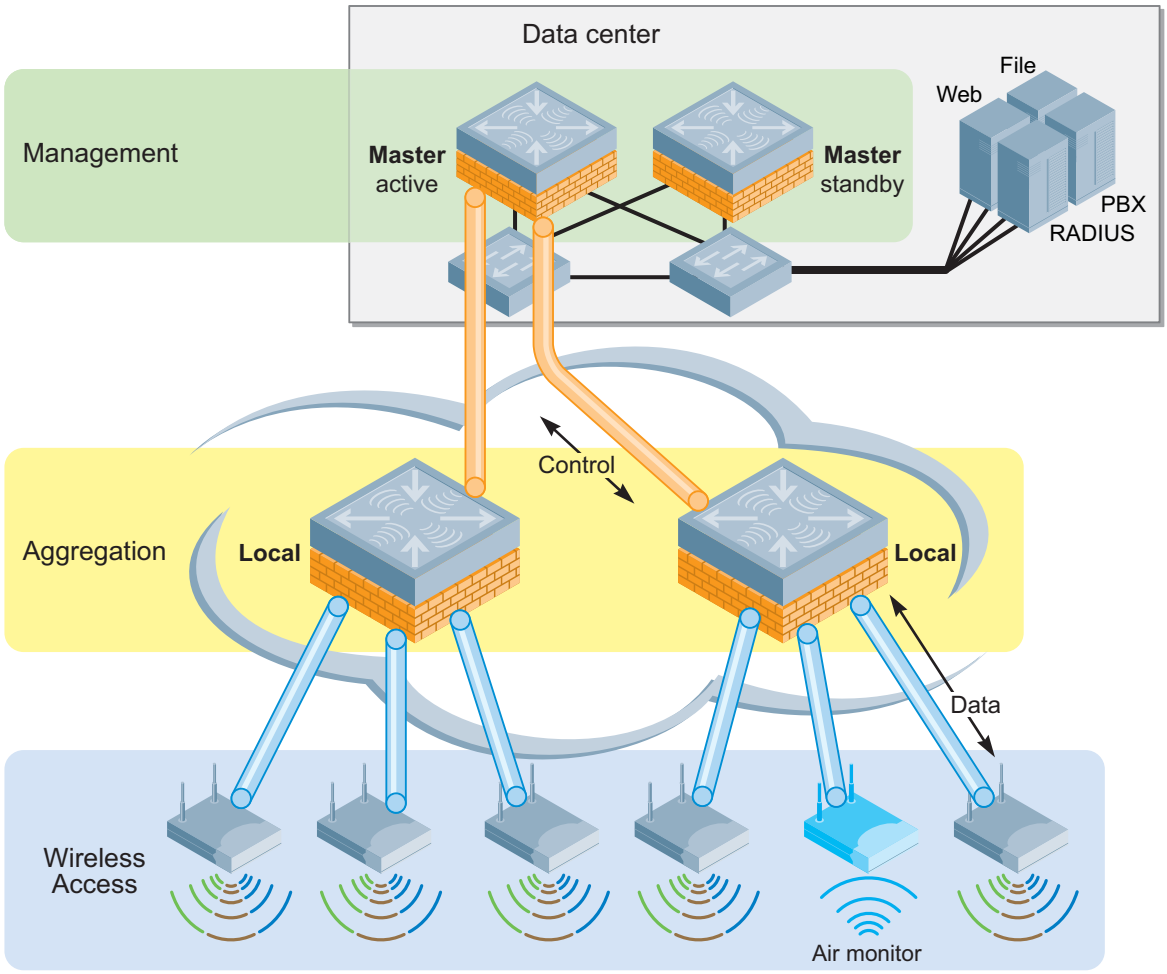

Together, the Management, Aggregation, and Wireless Access layers overlay on the Core, Distribution and Access infrastructure in a seamless, secure and high-performance manner. Any Aruba controller can serve as in the Management and Aggregation layer, and in smaller networks, a single controller provides both functions.

The network architect typically chooses the controller model that has capacity appropriate to the size of the user and AP population. In contrast to the Core/Distribution/Access model with capacity increasing as you approach the Core; a User-Centric network requires more capacity in the middle layer where tunnels are terminating and policies are being applied.

Other Aruba Reference Architectures

This Campus Wireless LAN Reference Architecture represents a large scale, highly available WLAN deployment model for a campus environment with numerous buildings that house thousands of users. This is the recommended deployment for this environment. There are other reference architectures that are considered best practices at different scales, and for different types of customer scenarios. Other Reference Architecture models that are commonly deployed by our customers are described in Appendix C on page 71.

Deployment of the Mobility Controller must be considered based on a number of factors; the most important of which is identifying where user traffic is ultimately destined. The Validated Reference Design for Campus Wireless Networks depicts the Master Controllers residing in the data center and Local Controllers deployed at the distribution layer.

Understanding Master and Local Operation

Once the controller count grows beyond a single pair of controllers, change control and network consistency can become an issue. To solve this management scalability issue, Mobility Controllers can be deployed in clusters consisting of a Master and one or more Local Controllers.

The Master Mobility Controller resides at the Management layer of the Aruba architecture in a data center environment. In an Aruba network employing a Master/Local design, configuration is performed on the Master and pushed down to the Locals. User troubleshooting, RF planning, and real-time RF visualization take place on the Master. The Master also controls Adaptive Radio Management (ARM) decisions for all Local controllers and is responsible for radio power and channel settings at the Wireless Access layer.

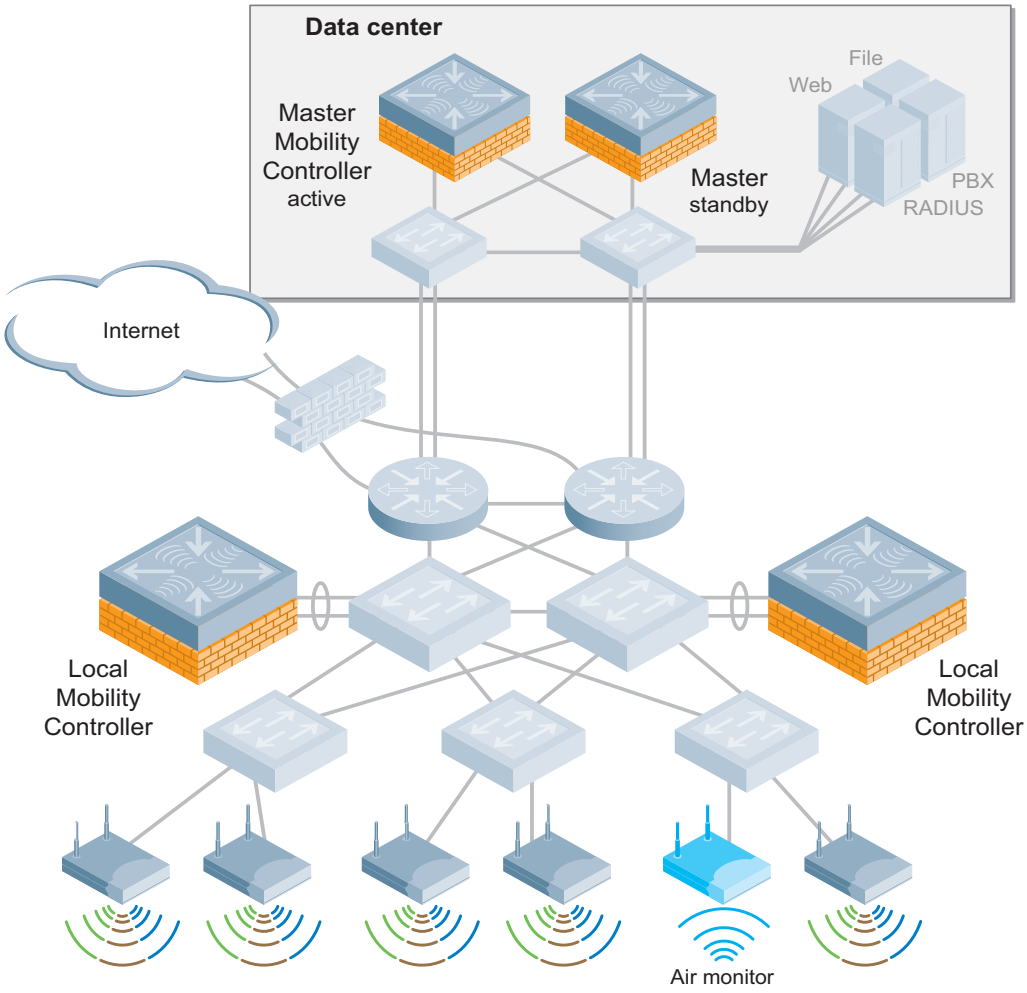

The Master is responsible for processing wireless intrusion detection system events, presenting the event and the corresponding wireless vulnerability and exploit (WVE) identifier. The Master is also responsible for handling location services correlation algorithms that compute the position of clients as well as rogue APs using signal strength measurements from APs in the network. All heat maps and location events will be handled through the Master Controller's web interface without needing an additional location appliance. This is the strategy depicted in the VRD model, and is the recommended model when two or more controllers exist in the same network.

In a large Campus WLAN with separate Management and Aggregation layers, Access Points and Air Monitors should never terminate on the Master Controller, they should only terminate on Local Controller.

If the Master becomes unreachable, the network will continue to operate as expected, but without the ability to perform operations such as configuration, heat map analysis or location services, until connection to the Master Controller is restored.

While the Master Controller is needed to perform configuration and reporting, it is not a single point of failure in the network..

Local Controllers reside at the Aggregation layer of the Aruba Overlay Architecture. They handle AP termination, user authentication, and policy enforcement. When configuring any Local Controller, you will need to know the IP address of the Master as well as the Pre-Shared Key used to encrypt communication between the controllers. If the Master becomes unavailable and no standby Master has been configured, the wireless network will continue to operate, but some management functionality will be unavailable until the connection is re-established.

The control channel between all Mobility Controllers is protected by an IPSec connection. This applies to both a data plane contained within the Local Controller, and a distributed control plane with some components on the Local and some on the Master Controller.

The controllers have a pre-configured key at first boot; this key must be changed for secure operation of the Master/Local cluster.

Mobility Controller High Availability

As Wi-Fi® networks move beyond conference rooms and become the primary network connection for users, the system must be robust enough to continue operation in the event of any network component failure. The Aruba system offers multiple configuration options to insure that the system operates in a highly available manner.

There are two different redundancies that must be considered: network management redundancy and network operations redundancy. Management redundancy is achieved by having redundant Master Controllers in the network at the Control layer; and operationally, by having two Local Controllers working together to share a load at the Aggregation layer, with each Local Controller acting as a backup for the other.

Master Controller Redundancy

To achieve high availability of the Master Controller, use the Master Redundancy method. In this scenario, two controllers are used at the Management layer with one controller configured as an active Master and one configured as a standby Master. The two controllers will synchronize databases and RF planning diagrams, and will run a Virtual Router Redundancy Protocol (VRRP) instance between them accessed by a Virtual IP (VIP) address. This is the address given to Access Points attempting to discover a Mobility Controller, and is used for network administration.

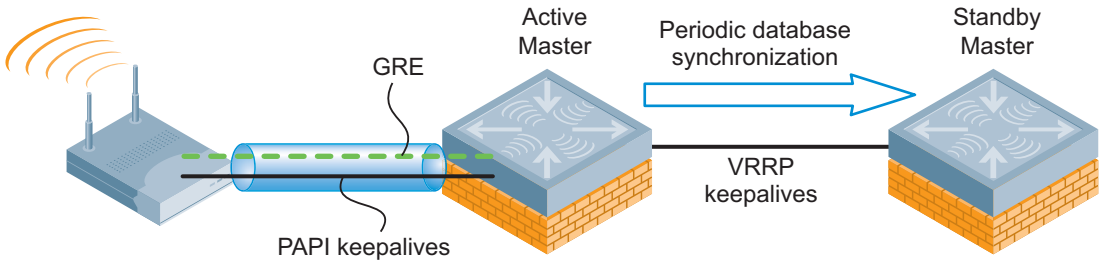

One Mobility Controller is always the Active Master Controller, and the other one is always the Standby Master Controller. Users managing the system will always log into the Active Master. It is not recommended that pre-emption be enabled on this setup. This configuration is known as "Active-Standby" redundancy.

In the Aruba Validated Reference Design, the recommended controller model to serve as a Master is the MMC-3600. The recommended network attachment method is to have each controller configured in a full mesh with redundant links to separate data center distribution switches.

Listed below is an example of the configuration of the “initially-preferred master”.

```
vrrp 22
vlan 22
      ip address 10.200.22.254
      priority 110
      authentication <password>
      description Preferred-Master
      tracking master-up-time 30 add 20
      no shutdown
database synchronize period 60
database synchronize rf-plan-data
```

The following shows the corresponding VRRP configuration for the peer Master Controller.

```
vrrp 22
vlan 22
      ip address 10.200.22.254
      priority 100
      authentication <password>
      description Backup-Master
      tracking master-up-time 30 add 20
      no shutdown
database synchronize period 60
database synchronize rf-plan data
```

Configure Local Controllers to use the VIP address as their Master Controller address as follows.

```
masterip 10.200.22.254 ipsec <key>
```

Local Controller Redundancy

Local Controllers at the Aggregation layer also use VRRP instances for redundancy, but in a different model than the Master Controllers at the Management layer. In this case, the controllers operate in what is known as “Active-Active” redundancy shown in the diagram below:

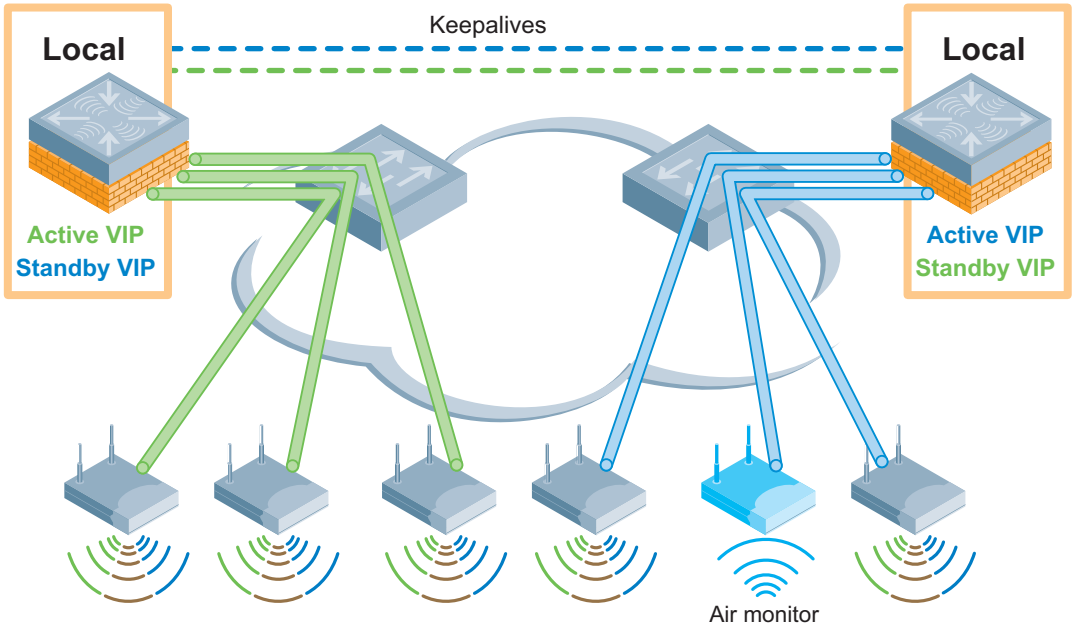

Using this model, two Local Controllers terminate APs on two separate VRRP Virtual IP (VIP) addresses. Each Mobility Controller is the active Local Controller for one VIP address and the standby Local Controller for the other VIP. The controllers each terminate 50% load of access points. The APs are configured in two different AP groups, each with a different VIP as the LMS IP address.

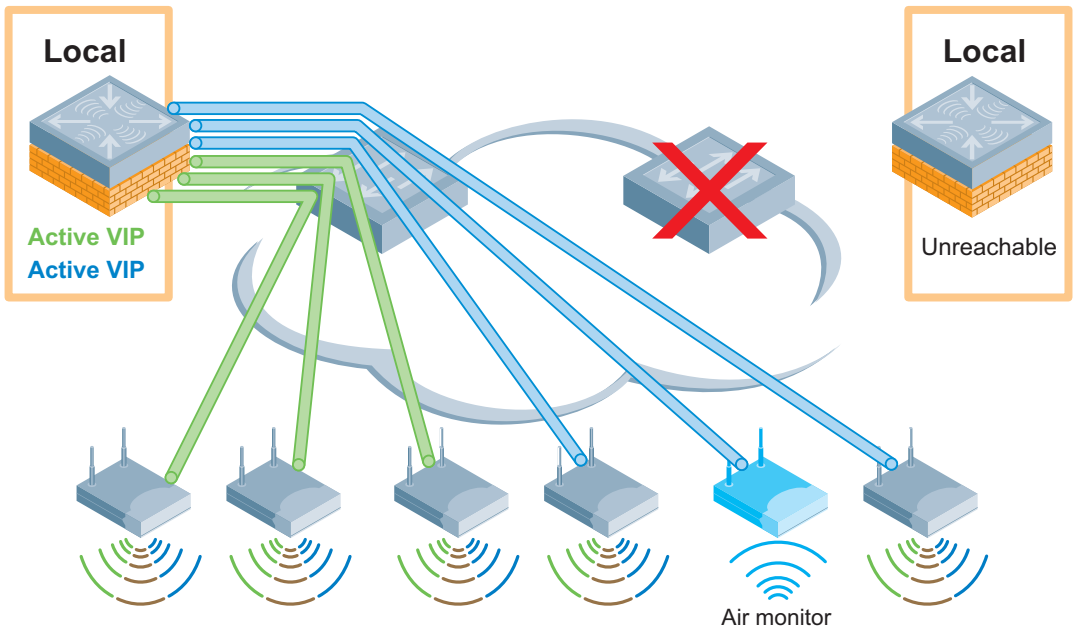

When one active Local Controller becomes unreachable, APs connected to the unreachable controller fail over to the standby Local Controller loading that controller to 100% capacity. Therefore each controller must have sufficient processing power and licenses to accommodate all of the APs served by the entire cluster. In this model, preemption should be enabled to force the APs to fail back to the original primary when it comes back online.

The configuration for each Local controller is a mirror image of the other. In the example below, the first controller is primary on 23 and standby on 24:

```
vrrp 23
vlan 23
      ip address 10.200.23.254
      priority 100
      preempt
      authentication <password>
      description initial-primary-23
      no shutdown

vrrp 24
vlan 24
      ip address 10.200.24.254
      priority 110
      preempt
      authentication <password>
      description initial-standby-24
      no shutdown
```

The second Local controller has an opposite configuration:

```
vrrp 24
vlan 24
      ip address 10.200.24.254
      priority 100
      preempt
      authentication <password>
      description initial-primary-24
      no shutdown

vrrp 23
vlan 23
      ip address 10.200.23.254
      priority 110
      preempt
      authentication <password>
      description initial-standby-23
      no shutdown
```

Using this scenario it is recommended to use the MMC-6000 chassis with redundant power supplies connected to at least two independent power sources. The recommended controller blade is the Multiservice Module. It is further recommended that these controllers have a “one-armed” connection to distribution layer switches, using Etherchannel to bond the two 10 Gigabit Ethernet connections.

N+1 designs are a common feature of other vendors’ centralized WLAN architectures. This is usually because the maximum number of APs that can be managed by one controller is limited to a few dozen or a few hundred at most, requiring the deployment of many controllers simply to service the

production AP load. By contrast Aruba supports up to 2,048 campus-connected APs and 8,192 Remote APs per controller which makes a 1:1 redundancy model feasible for the largest campus deployments.

With a properly implemented distribution layer, this Active-Active Local Controller design with VRRP at the Aggregation layer features full redundancy while offering performance advantages by load balancing during normal operation. This form of redundancy is superior to an N+1 design with a dedicated backup controller for the following three reasons.

- The network is already redundant – A properly implemented distribution layer has redundant links between access layer switches and core routers. If any link other than the ones to the Aruba Controllers fails, the system is already designed to route around that failure. Maintaining redundant links or having the Mobility Controllers 'straddle' between distribution layer switches does not add any additional reliability
- Loss of two controllers means a full network outage – Two Local controllers with physically isolated data connections on separate, redundant power sources are already protected against a majority of common failure modes. If both controllers lose power or link simultaneously it would most likely affect many more network components resulting in a complete network outage no matter how many redundant Local Controllers are available
- Better use of capital– In an N+1 design scenario at least one fully licensed Mobility Controller must always be sitting idle awaiting a network failure. Using Aruba's Active-Active capability allows both Local Controllers to terminate APs and enforce policies and user roles within the network, while providing hot backup for other members of the cluster

VLAN Design

When performing VLAN planning it helps to remember that VLANs are used in two logically different places on an Aruba Mobility Controller at the Aggregation layer. The first is the AP access side of the controller, where APs will terminate their GRE tunnels. These VLANs carry encrypted traffic back and forth between APs and the Controllers. The second is the user access side, where user VLANs will exist and where traffic to and from the user will flow. During authentication, a process called 'role derivation' assigns the proper VLAN to each user and forwards traffic to the wired network if allowed.

The user and access VLANs can also be visualized separately. In the first diagram below, the AP uses VLAN 100 for access. This represents the physical connection of the AP to the network.

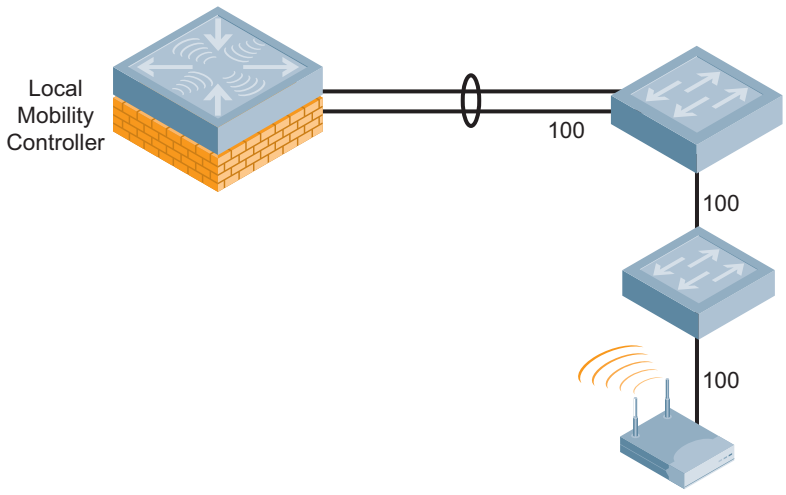

In the second diagram the client device is placed into VLAN 200 by the controller following completion of the role derivation process.

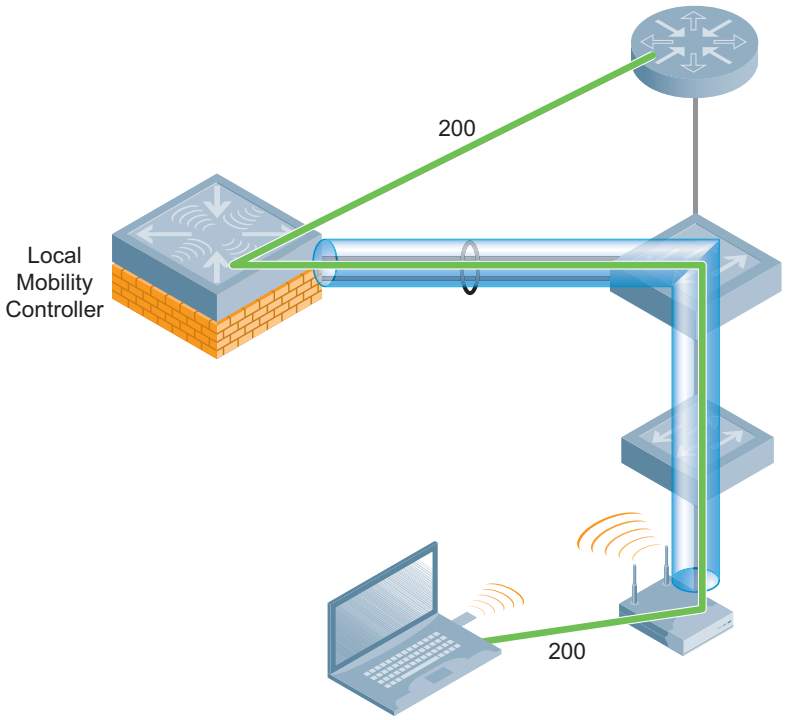

The user VLAN design will have implications for user connectivity and mobility across the network. To ensure that users do not overwhelm a single subnet, multiple VLANs can be configured to form a VLAN Pool in the Mobility Controller which users will be load balanced into dynamically. 'User mobility' is the ability of the user to roam between access points while remaining connected and not breaking user sessions through IP address changes.

Do Not Make Aruba the Default Router

The Mobility Controller is a Layer 3 switch that does not run routing protocols and should not be the default router for the VLANs on the network. The existing routers should remain the default gateways, with the Mobility Controller as a Layer 2 switched solution extending from the distribution layer.

Do Not Use Special VLANs

The use of 'special VLANs', which are VLANs created specifically for AP deployment, is not necessary and not recommended. No user traffic can enter the wired network except through the controller on which it terminates and after undergoing deep-packet inspection by the ArubaOS stateful firewall. As a result, there is no security risk to the network by putting APs on existing VLANs. In addition, for the Wireless Intrusion Detection System (WIDS) to operate properly, the Air Monitors need to see both the wireless and wired side of the network to properly classify rogue access points. When placed on isolated "AP VLANs", the WIDS system cannot correlate wired and wireless traffic. It will not be able to definitively classify rogue APs, and will not be able to automatically contain them.

VLAN Pools

Network administrators prefer to keep subnet sizes down to what is commonly referred to as class C network. This is a network with a subnet mask of /24 (255.255.255.0), yielding 253 user devices per subnet. This size is considered manageable and will 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.

One legacy methodology for dividing up large groups of wireless users is to take a set of APs and have all users associated with those APs placed in a single VLAN. Each set then would have one of the user VLANs associated with it. This method works if the user count never goes above the user count limit for that subnet; and if users have no need to roam outside of the AP group. 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.

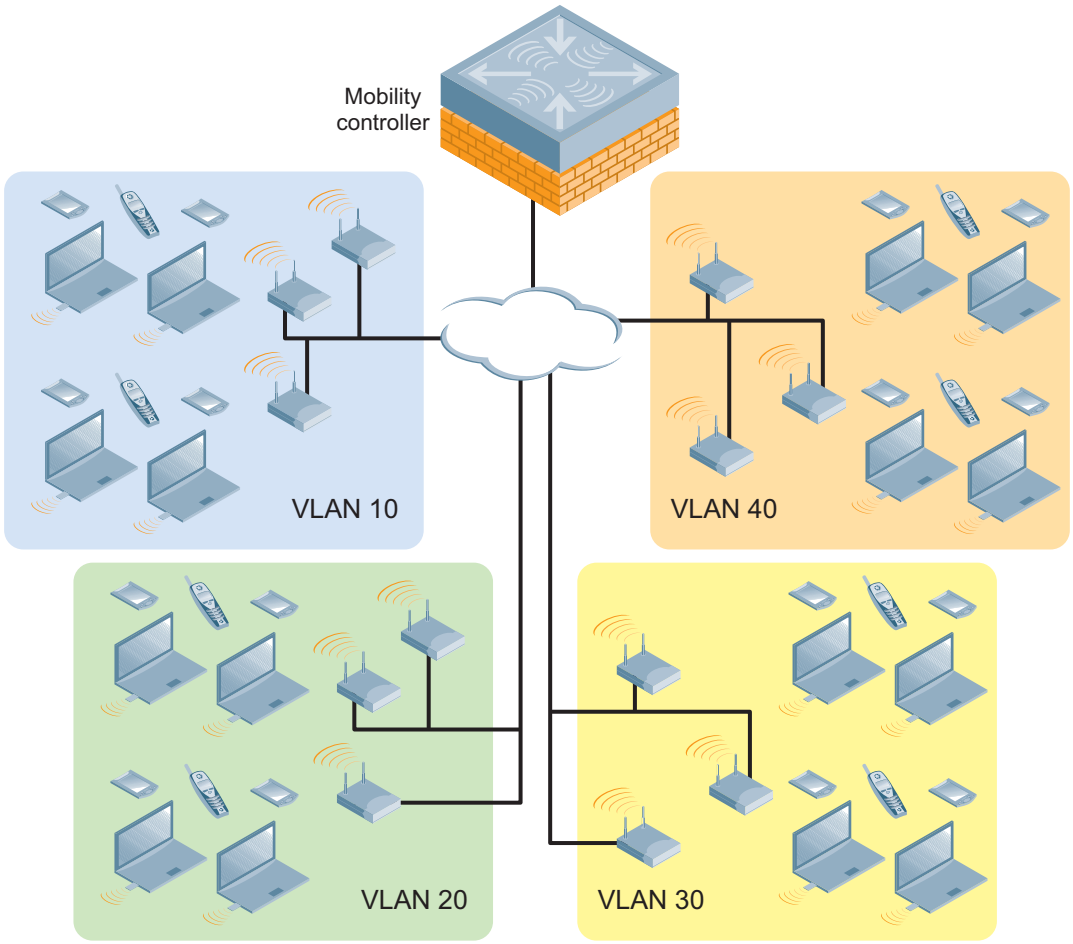

Aruba's VLAN Pooling feature allows a set of VLANs to be assigned to a designated group of users. These VLANs can be configured as a non-contiguous set, a contiguous range, or a combination of the two. As an example, the set could be VLAN numbers 10, 20, and 30. It could also be VLAN numbers 2 through 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 Aruba recommended method for handling user VLANs any time two or more user VLANs exist within the network.

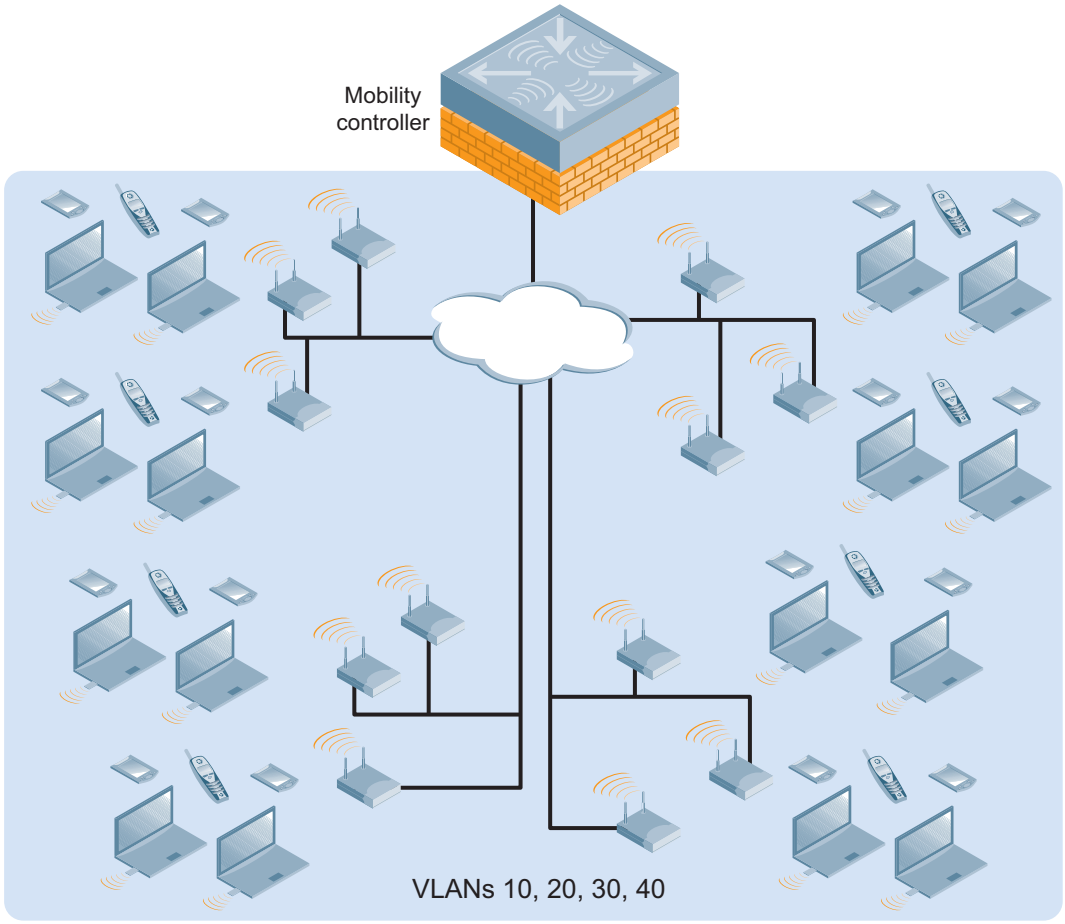

The system works by placing users in one of the VLANs in the pool. VLAN placement is done using the user's MAC address and running it through a hash algorithm. The output of this algorithm places the user into one of the VLANs in the pool and ensures that the user is always placed into the same pool during a roaming event. As the user associates with the next AP, their address is hashed. They are then placed into the same VLAN on the new AP and can continue to use their existing IP address with no break in their user sessions.

User Mobility and Mobility Domains

ArubaOS provides seamless wireless connectivity as users move throughout the network through its Mobile IP service. With roaming cutover times of 2-3 milliseconds, delay-sensitive and persistent applications such as voice and video experience uninterrupted performance. ArubaOS integrates proxy Mobile IP and proxy DHCP functions letting users roam between subnets, APs and controllers without special client software bringing IP mobility to any IP based Wi-Fi® device. The Mobile IP system is the only system that can scale once the system moves beyond a few controllers, as Layer 2 VLAN roaming will be far too cumbersome.

With Mobile IP, the ArubaOS will automatically tunnel traffic between a roaming client's original controller (the 'Home Agent') and the 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 controllers where their original subnet does not exist.

ArubaOS Mobility Domain

The ArubaOS Mobility Domain is the implementation of mobile IP addressing 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 re-authenticate 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.

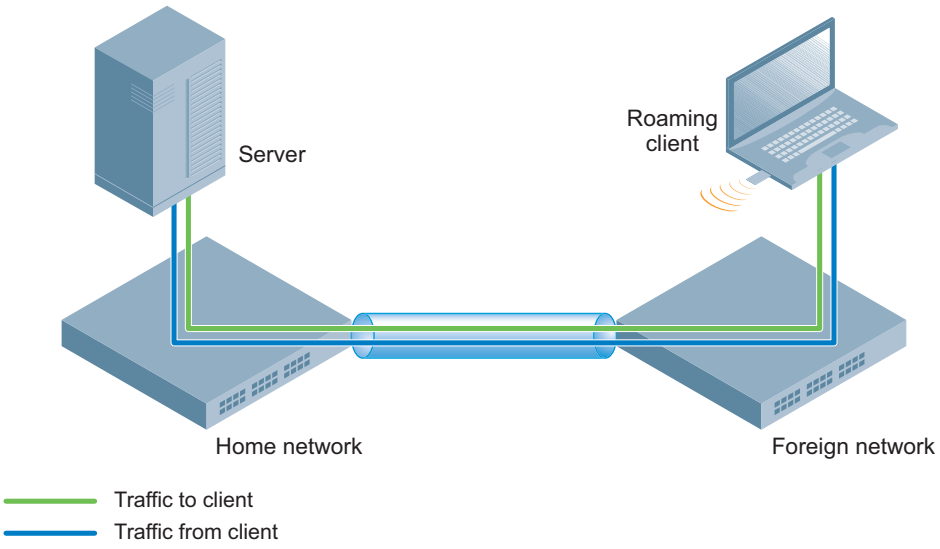

An Aruba Mobility Domain is a logical construct that defines a group of controllers physically close enough to one another that it could be reasonable that a user would roam between them in a single session. You can scale your Mobility Domain from a single domain on a limited number of controllers to multiple domains; each handling a separate country, campus, or building depending on your network design and business needs. 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. The Mobility Domain must be explicitly configured to allow roaming between the various controllers.

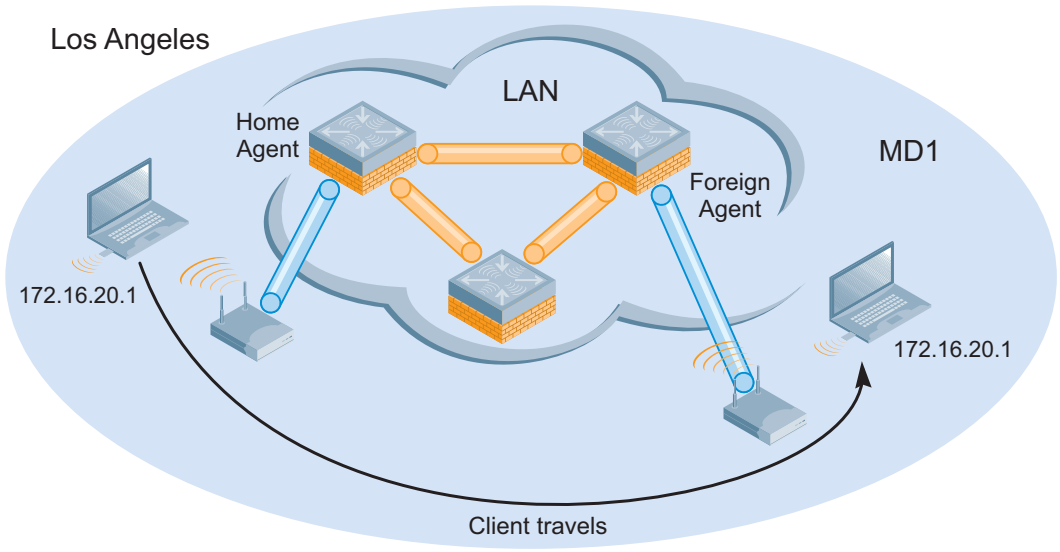

When the client roams off of its 'home' network to another network, the network is said to be attached to a 'foreign' network. The foreign network is defined as a network controlled by a different Mobility Controller than the one controlling the home network, but still within the same Mobility Domain. The IP address of the Mobility Controller on the foreign network becomes the client's 'care-of address'. This address is passed to the 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.

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

Mobility Domains take some amount of planning, but 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 roam between the subnets and controllers in a single session?" This 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 going to occur.

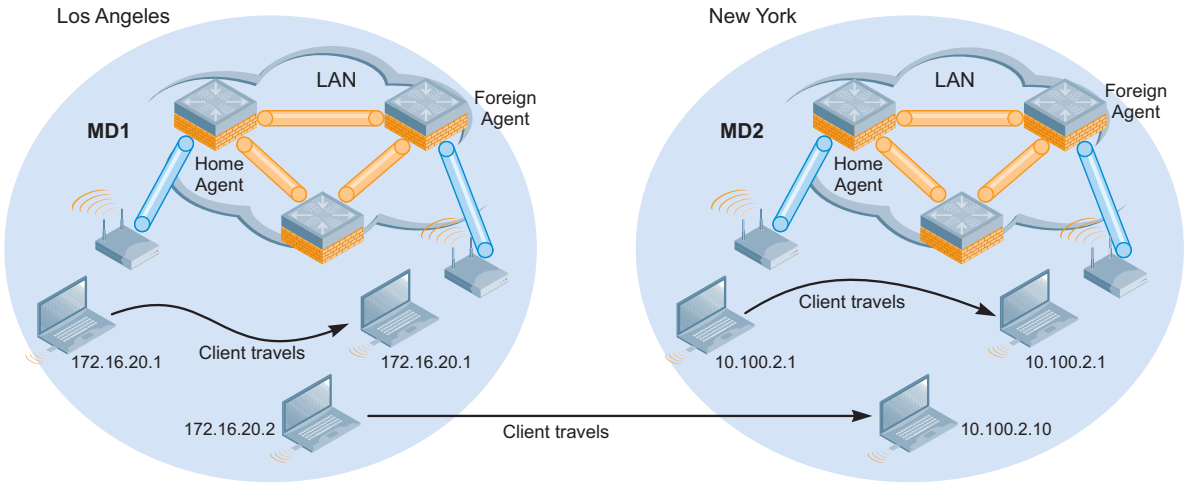

To plan a Mobility Domain, begin by taking a look at the network map, with a special focus on the access points and controllers. Generally, this will provide the information you need to develop a logical grouping of Mobility Domains. You should also examine heat maps of your network, and determine if the coverage areas provide enough connectivity and overlap to allow your clients to transition networks. Outdoor APs may extend this coverage between buildings providing you with a larger Mobility Domain.

Mobility Controller Physical Placement and Connectivity

Physical deployment of the Mobility Controllers is typically in two areas, the data center and the distribution layer of the network. The data center contains the Master Controllers that comprise the Management layer, while the distribution layer switches will connect to the Local Controllers that make up the Aggregation layer.

Master Controller Placement

The Master Controller should be given adequate bandwidth connections to the network, preferably a minimum of a Gigabit Ethernet LAN connection. Using the MMC-3600 appliance, Aruba recommends at

least two connections setting up redundant links to two data center distribution switches. With the Active-Standby configuration recommended in this VRD, this yields a full mesh between the two controllers and the distribution switches. The MMC-3600 does not have redundant power supplies; connect each appliance to discrete power sources in the data center.

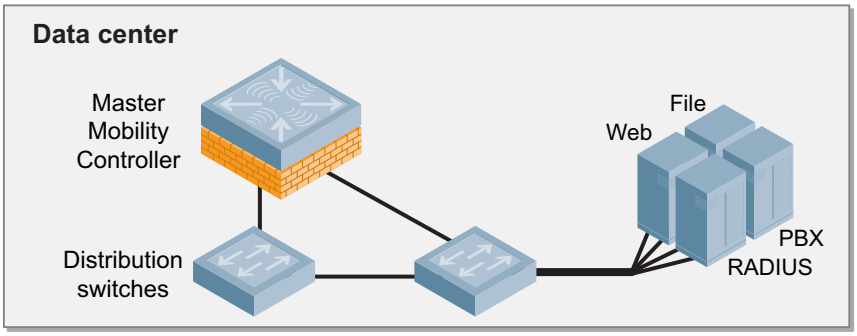

Local Controller Placement

The Local Controller should be connected to the distribution layer switches in an MDF or similar suitable location with backup power, with each Active-Active pair connecting to separate switches. Using the MMC-6000 Multiservice Mobility Controller as recommended by this VRD, each blade should be connected to its own distribution layer switch with two 10 Gigabit Ethernet connections bonded with Etherchannel. A fully populated MMC-6000 chassis with four blades require eight Gigabit connections. Each MMC-6000 chassis should contain redundant power supplies connected to discrete power sources.

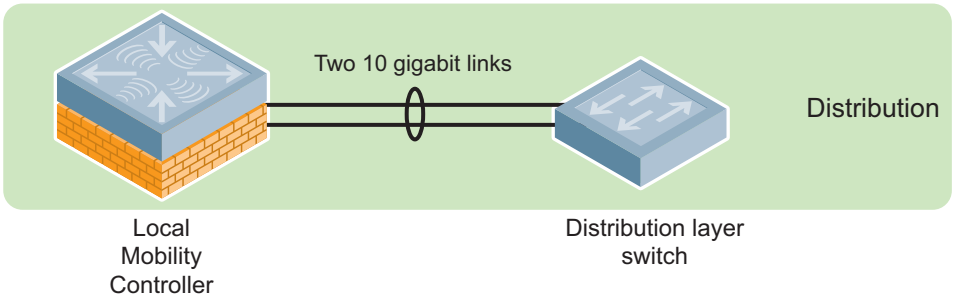

AP Placement, Power, and Connectivity

Mobility Controller and Thin AP Communication

Mobility Controllers and thin APs work as a system. Configuration for all APs is automated and centralized on the Mobility Controller. Upon bootup, each AP uses DHCP to obtain its IP information then connects to the Local Controller to retrieve its initial configuration, and to check for firmware updates. Subsequent configuration changes are performed centrally within the Local Controller and pushed to each AP. If the firmware on the AP does not match the controller, the AP will automatically use either FTP or TFTP to upgrade itself to the new firmware stored on the Local Controller with no administrator intervention.

Communication between the AP and the Local Controller at the Aggregation layer occurs using a GRE tunnel established during the boot process. Because the GRE tunnel is in place, all wireless traffic is transmitted directly to the controller, so no special VLANs need to be deployed for APs; they will function over the existing infrastructure as would any other client. This avoids the “VLAN explosion” problem in some other architectures where every user VLAN must terminate on every AP throughout the enterprise. On the other side of the GRE tunnel, the user traffic is then switched to the correct

VLAN at the Local Controller where a VLAN trunk already exists. This also allows for mobile IP functionality without client software as the intervening VLAN between the AP and the controller is never seen by the client.

AP Power and Connectivity

The AP can use DHCP for IP addressing and can automatically discover the Mobility Controller through a number of methods making it easily added to any existing employee port and VLAN.

If the Access Point and Mobility Controller share the same Layer 2 VLAN, then nothing else needs to be done as the AP will automatically discover the controller via the Aruba Discovery Protocol (ADP). If the AP and controller are separated by a Layer 3 network then two other methods are available for controller discovery. An entry can be entered into the organization's DNS for 'aruba-master' with the AP address of the Mobility Controller, or a DHCP option 43 value may be configured with the address of the Mobility Controller.

Power should be supplied either using 802.3af Power-over-Ethernet (PoE) or using a power adapter for the access point. PoE is the simplest method if it is already in place because the AP will be able to use a single cable for both power and data.

AP Location and Density Considerations

Determining the correct number of APs to deploy for a given area requires careful planning. RF designers generally use a metric called 'AP density' which refers to the number of square feet that each AP is expected to serve.

AP Density is affected by:

- What applications are running on the network
- Which frequency bands are in use
- The degree of overlap in the cells for self-healing
- Antenna choice

It is possible for AP density to vary within a campus or even within a building. Aruba recommends working with a professional WLAN engineering organization to select the proper AP density for all coverage areas.

In addition to AP density, the RF engineer must also select a Placement Methodology. This refers to whether the APs are spaced uniformly or not, and whether they are located along the perimeter of an area or spread throughout the interior. The methodology has important consequences for customers that plan to use location services. With the AP Density and Placement Methodology known, the RF engineer can use the Aruba RF Plan tool to create a design for each floor or area to be covered. This is explained in more detail in a later chapter.

Office Deployment

It used to be common practice to deploy APs in the plenum space above the ceiling grid. As the cost of APs came down, and with the advent of the thin AP with attractive industrial design, it is now common to deploy the AP directly in the user space. Access Points such as the Aruba AP65 have an integral T-bar clip to attach directly to standard drop ceilings common in carpeted office areas, as well as supporting many other mounting methods. A best practice is to clearly label the AP during installation using letters large enough to be seen from the ground to facilitate moves, adds and changes.

Performance is best when a clear line-of-sight (LOS) exists between the AP and its clients. Aruba does not recommend placing the AP on desktops, or placing an AP on the top of a set of cubicles. LOS is easily obstructed in these cases, resulting in performance that may not meet the standards of the design.

Be sure to remember that RF travels in three dimensions. In a multi-floor building, the strongest signal may be above or below rather than side-to-side. In all 802.11 networks the client, rather than the AP, makes the decision when to roam from one AP to another. The RF designer can use this to advantage by staggering APs from floor-to-floor. This will help ensure that client roaming behavior is predictable, and can improve how ARM makes decisions about channel selection and power settings.

Voice Deployment

When using Voice over WLAN take into consideration the type of handset that will be used. Many older voice handsets are only capable of operating in the 802.11b frequency range (2.4Ghz). To provide the highest quality of service, Aruba recommends moving clients into the 802.11a band due to the greater number of available channels, and using a higher AP density. This approach requires dual-radio APs to service both client types, and dual-radio AMs to lock the air in both bands.

The cell design and AP density is also affected by handset manufacturers. Generally speaking, a voice network should be RF planned to provide a minimum signal strength of -67dBm or better throughout the service area. In the Aruba RF Plan tool, use a 150% overlap setting with a 54 Mbps minimum data rate to provide this level of coverage. In most cases, this translates to an AP approximately every 60 feet.

Active RFID Tag Deployment

Placement Methodology and AP Density are both important when using active RFID tags. Because Aruba's RF Locate feature uses triangulation of Received Signal Strength Indication (RSSI) to locate devices and active RFID tags, the tag or device must be heard by a minimum of three access points to obtain a reliable reading. APs should be deployed along the building perimeter so that there is always a defined edge the client or tag will be contained within. If necessary, external semi-directional antennas can be used on these perimeter APs to direct the maximum signal towards clients and to reduce susceptibility of the system to co-channel interference from outside the building.

Chapter 6

Mobility Controller Configuration

Once the hardware has been deployed there are several design decisions required to build out a working production network. This includes VLAN and IP network design, as well as the loop back IP address selection and spanning tree usage. Many of the decisions will logically follow from where the network architect chooses to place the AP and controller in relation to one another.

Other items needing additional planning are:

- Configuration Profiles and AP Groups
- Service Set Identifier (SSID) selection
- Authentication and encryption methods

This chapter will provide a brief introduction of these topics with additional detail provided later in the document.

Required Licenses

To build this Aruba Validated Reference Design for a large Campus as described in Chapter 4 on page 19, the following licenses are required on the Local Controllers, assuming an MMC-6000 Multiservice Mobility Controller is acting as a backup to a second MMC-6000:

- LIC-512-AP Access Point License (512 Access Point License)
- LIC-WIP-512 Wireless Intrusion Protection Module License (512 AP Support)
- LIC-PEF-8192 Policy Enforcement Firewall Module License (8192 Users)
- LIC-VOC-8192 Voice Services Module License (8192 Users)

The following licenses should be applied to the Master Controllers assuming a MMC-3600 controller with no APs terminating and not acting as a backup for any active controller:

- LIC-8-AP Access Point License (8 Access Point License)
- LIC-WIP-8 Wireless Intrusion Protection Module License (8 AP Support)
- LIC-PEF-128 Policy Enforcement Firewall Module License (128 Users)
- LIC-VOC-128 Voice Services Module License (128 Users)

Configuration Profiles and AP Groups

Configuration profiles and AP Groups work together to provide an abstraction layer between the physical settings of the system, and the conceptual goals of the network architect. This abstraction feature provides the Aruba administrator with the benefits of reusable groups of settings (called 'profiles') that can be applied in a mix-and-match fashion with extremely fine granularity.

Configuration Profiles

Configuration Profiles allow different aspects of the Aruba system to be grouped into different configuration 'sets'. Each profile is essentially a container, and the container creates a particular configuration based on settings within the container. SSID Profiles, Radio Profiles and AAA Profiles are just some of the available choices; and each one includes a number of parameters that can be adjusted to meet the needs of the design. Multiple versions of the same profile can be created and given different

names. This allows the administrator to define a particular profile once and reuse it as needed which reduces errors and data entry.

The ArubaOS Profile system is set up so that the configuration flow goes from high level to low level in a hierarchical manner. Unlike other hierarchical systems such as LDAP, the system does not provide arbitrary levels of depth or inheritance. The ability to copy a profile when creating a new profile allows for rudimentary inheritance when the new profile is created. Changes to the original profile will not be reflected in the new profile.

Profile Types

The basic idea of a profile is very straightforward. With nearly 30 types of profiles available, ArubaOS 3.3 offers the administrator almost unlimited control over how their wireless network can be implemented. The main categories of profiles are shown below. Each box represents a different profile. Note that certain profiles are nested within others.

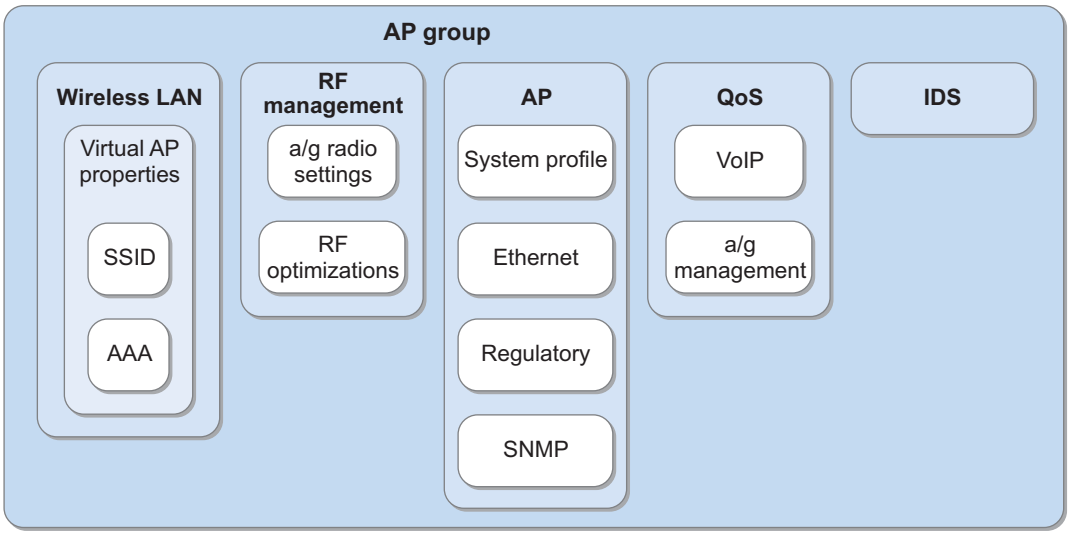

Some of the more common profiles administrators work with daily include:

- AP Profiles

  Configure AP operation parameters, radio settings, port operations, regulatory domain, and SNMP information.

- QoS Profiles

  Configure traffic management and VOIP functions.

- RF Management Profiles

  Configure radio tuning and calibration, AP load balancing, coverage hole detection, and RSSI metrics.

- IDS Profiles

  IDS functions for APs. There is a top-level IDS profile that contains other IDS profiles in which you configure detection of denial of service (DoS) and impersonation attacks; unauthorized devices on the wireless network, as well as intrustion signatures.

AP Groups

An AP Group is a unique combination of Configuration Profiles. In general, all profiles are available to be assigned to an AP Group to create a complete configuration. This flexibility in configuration allows you to do arbitrary groupings of APs such as 'All Lobby APs' or 'All APs in California' with different configurations on each. Each AP Group must include a minimum number of profiles, in particular, a Virtual AP Profile.

It is important to note that each Access Point or Air Monitor can be a member of only a single AP Group. You can not assign multiple AP Groups to the same AP. This restriction prevents the assignment of incompatible or conflicting Profiles.

Profile Planning

To effectively use the profile system takes some planning. Unlike most planning decisions in network designs, profile planning is not based on performance and scalability; it is based on creating a functional and flexible network design that can be logically understood. Ideally, this planning is part of the network planning.

While it is possible to simply place all of your equipment in default profiles and change the parameters to suit your needs, you will miss out on the power and flexibility of the system. To take full advantage of the system you must take into account the physical layout of your equipment, the technical management requirements, and the business practices and regulatory requirements specific to your organization.

Aruba recommends changing the following defaults:

- Default AP-Group
- Default Virtual-AP
- Default SSID.

When an AP first boots, it is automatically made a member of the default AP-Group, which has open authentication by default. Aruba recommends changing the default to Air Monitor mode for new Access Points. This allows anyone who plugs an unauthorized Aruba AP into your network to simply add to your monitoring capabilities instead of creating a backdoor.

SSIDs, VLANs and Role Derivation

Each Aruba Access Point has the ability to appear to wireless users as multiple physical APs. Each of these 'virtual APs' has their own Basic Service Set Identifier (BSSID) that identifies the AP and the network name, or Service Set Identifier (SSID).

SSIDs

SSIDs appear as the name of the network displayed in the 'Available Wireless Networks' screen on a wireless client. While many APs in the same network will share the same SSID, each will have a unique BSSID. This feature is often used to let users know which SSID they should attempt to associate to, and to provide different levels of security to each of the SSIDs, such as WPA, WPA2, and Captive Portal. Clients typically make roaming decisions based on the received signal strength of the audible BSSIDs they can hear.

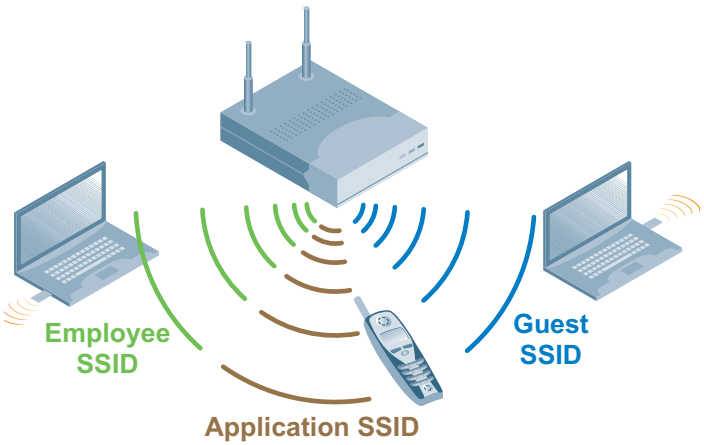

The diagram above shows the most common SSID design for enterprise organizations that includes three different SSIDs. A strong authentication and encryption suite is used for employee users, in this case WPA2 - Enterprise. The network administrator might choose a name something like 'Acme Corp Employee' for this SSID.

The second SSID is used for specific devices which are not capable of modern high authentication and encryption levels. As of this writing, common examples includes the following devices:

- Portable barcode scanners
- Active RFID tags
- All but the latest WiFi phones
- IP video cameras

In this case, the Mobility Controller uses an SSID such as 'Acme Corp-Application' and uses the strongest authentication and encryption suite supported by the devices; in this case, WPA-PSK (pre-shared key).

The final SSID is used to provide guest access to the network. This SSID will not run any encryption and will require guests to authenticate using the Captive Portal capability that is built into the Aruba Mobility Controller. The guest users can authenticate against a centralized authentication server or the built-in Local Database on the Mobility Controller; which is common when combined with the guest provisioning role on the controller.

VLANs

At the controller, users who successfully authenticates via an Aruba AP into any of these three SSIDs are treated very differently in the Role Derivation process according to the Configuration Profiles in the AP Group assigned to that AP. The Employee user is most likely placed on a VLAN with access to internal network resources, although this can be further refined with sophisticated ACLs applied on a per-packet basis. The dual-mode WiFi phone is placed on a voice-only VLAN and only permitted to contact a SIP server and transmit RTP traffic. Any attempt by the device to do something else would automatically 'blacklist' that device from the network. Finally, the Guest user would be placed onto a guest-only VLAN that only has access to the default gateway leading to the internet.

Role Derivation

Aruba uses the term 'Role Derivation' to describe the process of determining which role is to be assigned to a user. The system can take into account the user's credentials, location, time of day, and authentication type when deciding which role to assign.

This system can be as detailed or as general as the administrator prefers. The Role Derivation process determines:

- What class of service is provided to user traffic
- Which Firewall ACLs are applied to the user's traffic
- Which VLAN the user is placed into

Secure Authentication Methods

The most common authentication methods for Campus WLANs are 802.1X, and Captive Portal; other authentication methods are also discussed in this section. Mobility Controllers at the Aggregation Layer are the central point of control for users and access points, and are typically deployed in the distribution layer of the network. The Mobility Controllers sit in the authentication path, terminate user-encrypted traffic, and enforce policy using the optional Aruba Policy Enforcement Firewall module.

This ICSA certified stateful firewall allows control of user traffic as well as application awareness through deep packet inspection. The Aruba Policy Enforcement Firewall module has the capacity to dynamically follow sessions, log user sessions, and take actions through the blocking of user traffic and blacklisting of users for policy violation. This Role-Based Access Control system allows users with different access rights to share the same access points.

A wireless user gains access to the network by attempting to associate to the AP with the strongest signal. The association request may have originated from a new user logging on to the network, or an active user who has just roamed to a different location. The 802.11 MAC layer protocol association request is forwarded to the Mobility Controller, which then attempts to retrieve the user's state from the active user database. If the user was not active previously, the Mobility Controller will proceed to authenticate the user using 802.1X coupled with back-end authentications mechanisms such as RADIUS, Active Directory or LDAP.

The Mobility Controller can perform user authentication in multiple ways to suit the varying needs of an enterprise, and the existing AAA infrastructure in use. The most typical authentication methods employed on Aruba networks can be summarized as:

- 802.1X based user authentication with a backend server
- 802.1X PEAP termination on the controller
- PPP based user authentication over IPSec based VPNs
- Captive Portal based user authentication
- A combination of authentication methods such as 802.1X followed by captive portal, or WEP authentication followed by VPN

Authentication in the Aruba system typically leverages existing authentication stores, including RADIUS, Active Directory, and LDAP. While the Aruba Mobility Controller does contain a scalable Local DB for users and guests, it is typically desirable to have that functionality leveraged from an existing authentication system to ease synchronization issues.

Authenticating with 802.1X

802.1X was developed to secure wired ports by placing the port in a 'blocking' state until authentication completed using Extensible Authentication Protocol (EAP). EAP is a framework and allows many different authentication types to take place within the EAP authentication system; Protected EAP (PEAP) is most commonly used in wireless. In this mode, a TLS tunnel is created and user credentials are passed to the authentication server within the tunnel. When the authentication is complete, the client and the Mobility Controller both have copies of the keys used to protect the user session.

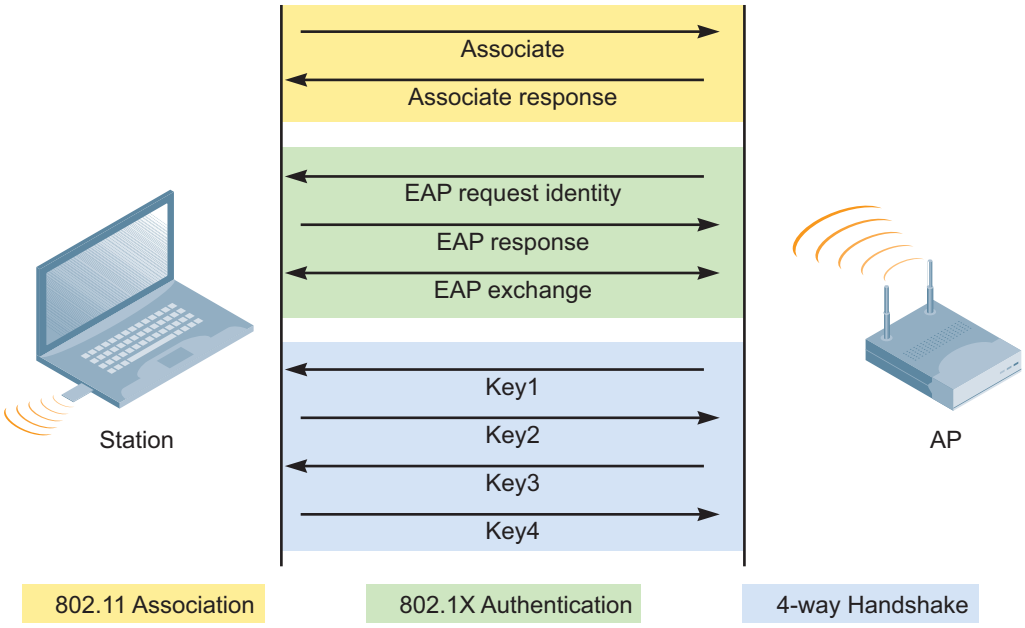

Using RADIUS and a WPA2 protected connection as an example, authentication occurs using 802.1X. The Mobility Controller forwards the request to the RADIUS server who performs the actual authentication and sends a response to the Mobility Controller. Once authentication completes successfully, encryption keys are passed to the Mobility Controller from the RADIUS server, along with the user's access policies. The Mobility Controller then completes the role derivation process and adds the new user, along with all the relevant state information, into the active user database and completes the authentication process. A security context is created, and for encrypted links, key exchange occurs where all traffic is now encrypted.

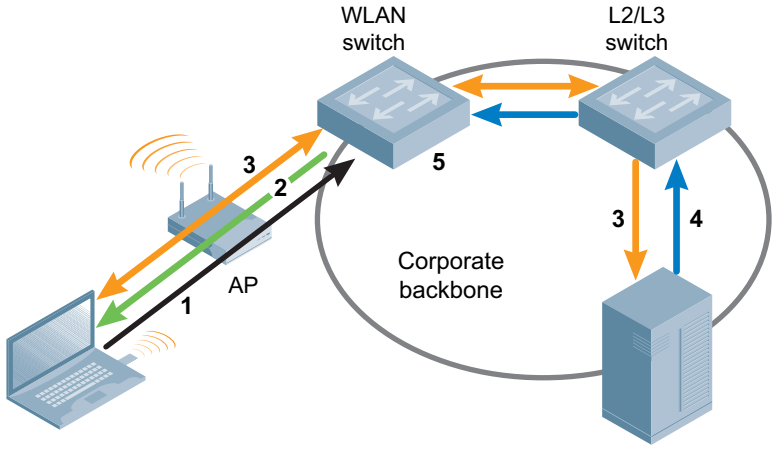

1. Client sends 802.11 association request that is automatically forwarded by the AP to the WLAN switch
2. WLAN switch responds with association acknowledgement
3. Client and WLAN switch start 802.1X authentication conversation along with RADIUS server
4. Encryption keys passed to the WLAN switch, and user derives own encryption keys, begins sending encrypted data
5. WLAN switch decrypts data, processes packets, applies services and forwards packets based on .11 MAC

If the user already exists in the active user database and is now attempting to associate to a new AP, the Mobility Controller will understand that an active user has moved, will restore the user's connectivity state and initiate mobility processing.

ArubaOS uniquely supports AAA FastConnect™, which allows the encrypted portions of 802.1x authentication exchanges to be terminated on the Mobility Controller where Aruba's hardware encryption engine dramatically increases scalability and performance. Supported for PEAP-MSCHAPv2, PEAP-GTC, and EAP-TLS, AAA FastConnect™ removes the requirement for external authentication servers to be 802.1x-capable and increases authentication server scalability by permitting several hundreds of authentication requests per second to be processed.

Authenticating with Captive Portal

For clients that do not support WPA, VPN, or other security software, Aruba supports a Web-based captive portal that provides secure browser-based authentication. Captive portal authentication is encrypted using SSL (Secure Sockets Layer), and can support both registered users with a login and password or guest users who supply only an email address. Through Aruba's integrated Guest Connect system, captive portal can provide a secure guest access solution by permitting front-desk reception staff to issue and track temporary authentication credentials for individual visitors.

The user connects to the SSID, which requires no authentication, and is placed in a state that requires a login. When the user opens a web browser they will be presented with a captive portal screen asking them to enter credentials, enter an email address, or simply accept a set of service terms.

Authentication Methods for Legacy Devices

Other authentication methods include pre-shared keys (PSK), Wired Equivalent Privacy (WEP), and open access with no authentication or encryption. Pre-shared keys are often used on older devices, or devices which cannot handle full 802.1X authentication. WEP is typically only used by very old devices, and due to the ease of which the system can be compromised should not be used at all if possible. Open authentication is typically used in open hot spots where there is no requirement to authenticate or secure customer usage and is rarely used.

Configuring Roles for Employee, Guest and Application Users

The Aruba system is unique; it combines user-based security as a part of the WLAN model. When a user is authenticated, using one of the methods discussed in the previous section, a role is applied to the user that is enforced via the firewall and the defined policies for that user.

Employee Role

Users who are company employees can be granted a role based on their specific job function, or simply be given a universal 'employee role'. Additional granularity can be applied, such as permitting a user in engineering to access the engineering subnets but not the finance or accounting servers.

In smaller organizations, users will most likely be placed in a single user subnet that has access to all internal and external resources.

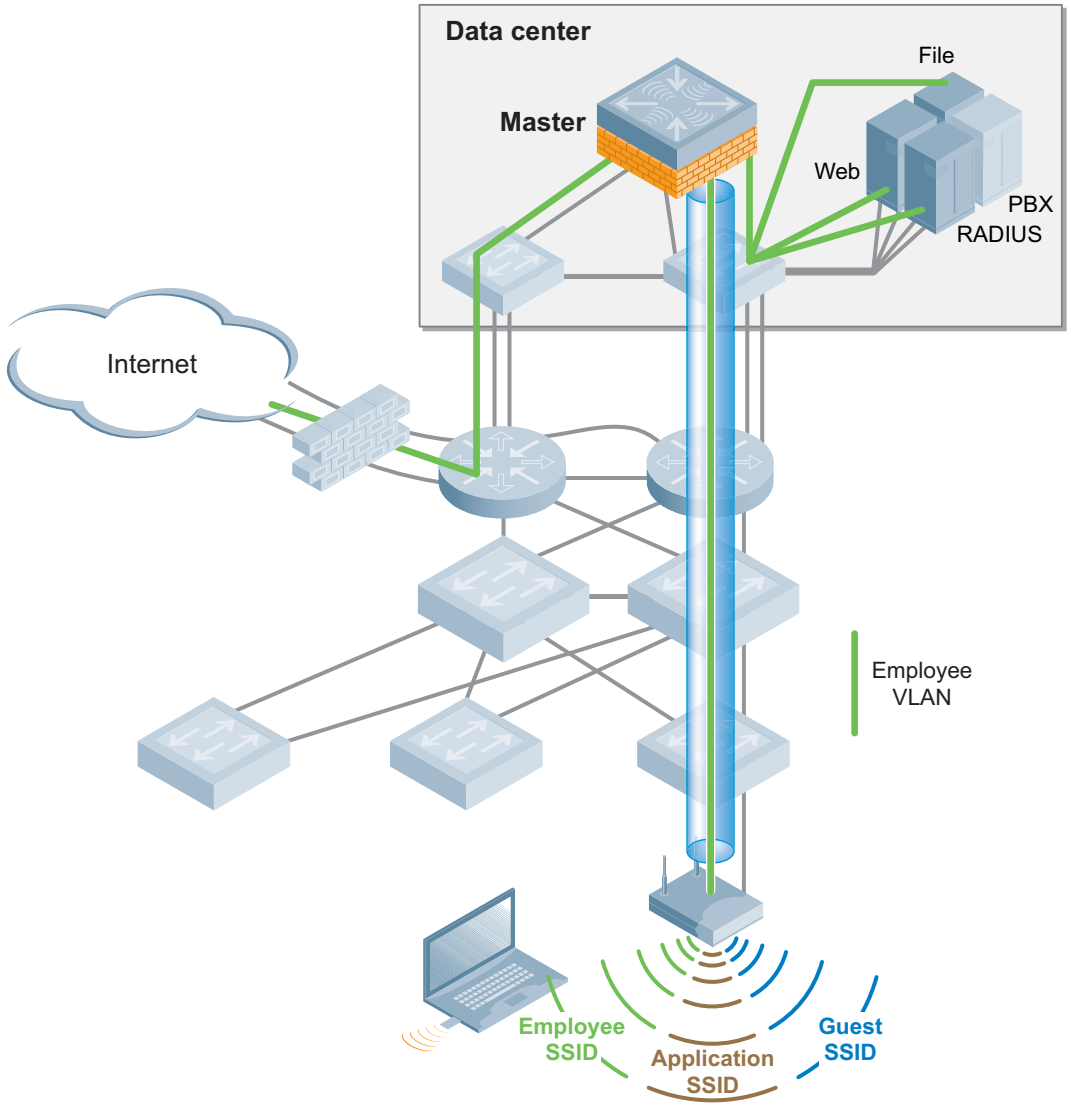

Guest Role

Guest usage warrants special consideration for enterprise wireless networks. It is not enough for guest users to be separated from employee users through VLANs in the network. Guests must be limited not only in where they may go, but also limited by what network protocols and ports they may use to access resources.

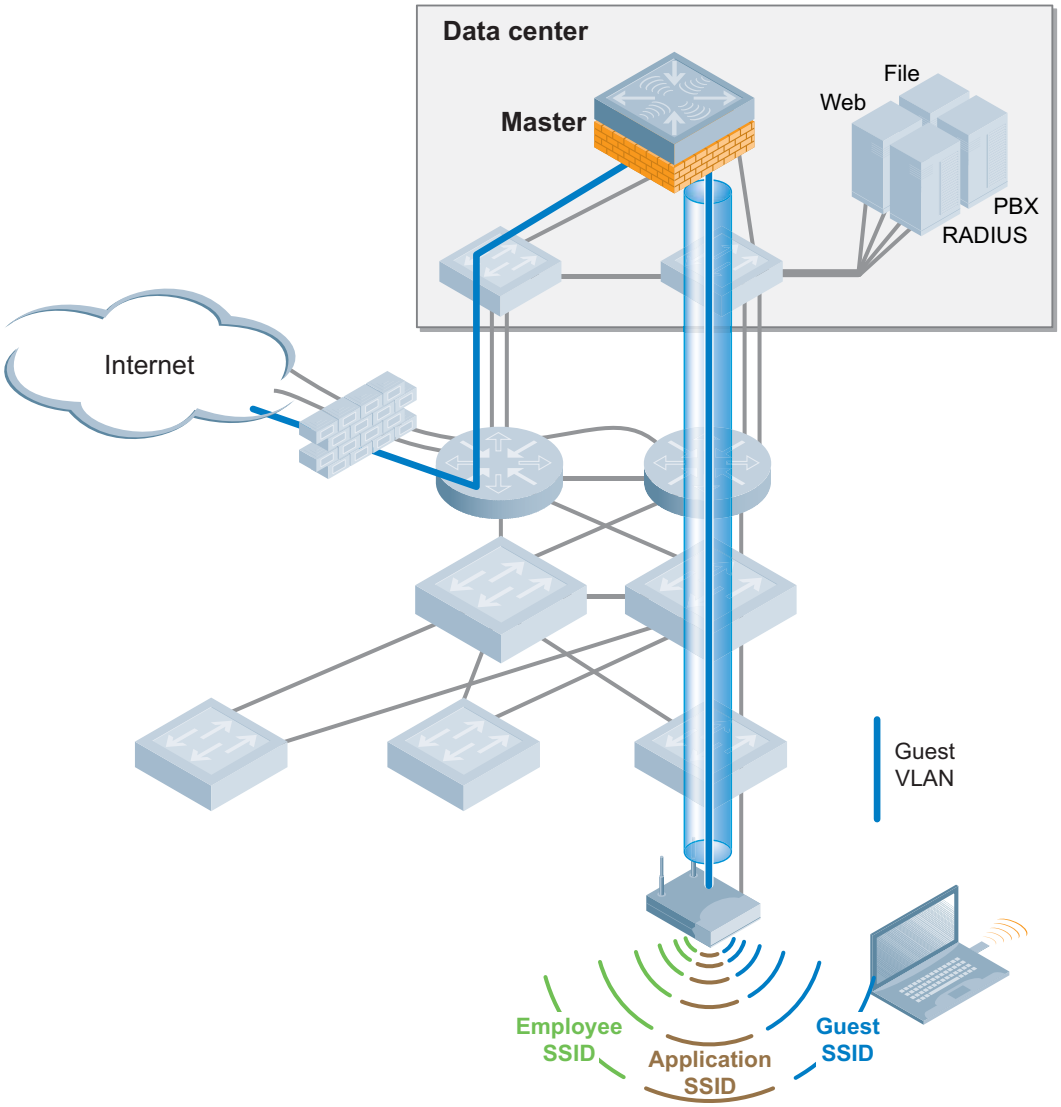

Good guest policy as implemented by the stateful firewall should only allow the guest to access the local resources that are required for IP connectivity. These include DHCP and possibly DNS if an outside DNS server is not available. All other internal resources should be off limits for the guest. This is usually achieved by denying any internal address space to the guest user.

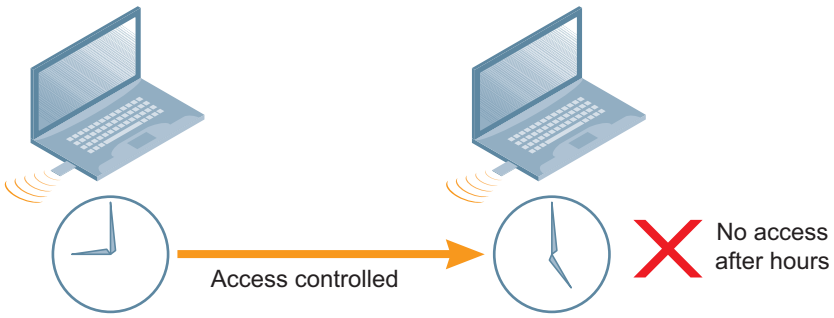

Additional policies should be put in place to limit the use of the network for guests. The first policy is a time-of-day restriction. The user should be limited to accessing the network during normal working hours as they should only be using the network while conducting official business. Accounts should be set to expire when their Local work is completed, typically at the end of each business day.

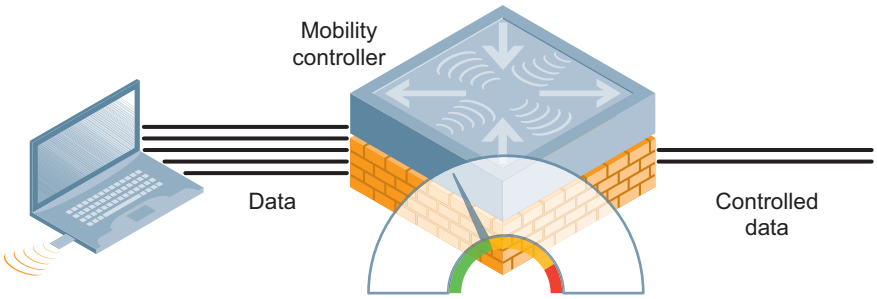

A rate limit can be put on each guest user to keep the user from using up the limited wireless bandwidth. Employee users should always have first priority to the wireless medium for conducting company business. Remember to leave enough bandwidth to keep the system usable by guests. Aruba recommends a minimum of 10%. Guests can always burst when the medium is idle.

Create a time range:

```
time-range working-hours periodic
   weekday 07:30 to 17:00
```

Create a bandwidth contract and apply it to an AP group:

```
wlan traffic-management-profile "employee-guest-app"
   bw-alloc virtual-ap “corp-employee” share 45
   bw-alloc virtual-ap "corp-app" share 45
   bw-alloc virtual-ap "guest-net" share 10
ap-group "corp-aps"
   dot11a-traffic-mgmt-profile " employee-guest-app"
```

Create aliases:

```
netdestination “Internal-Network”
   network 10.0.0.0 255.0.0.0
   network 172.16.0.0 255.255.0.0
   network 192.168.0.0 255.255.0.0
netdestination “Public-DNS”
   host 64.151.103.120
   host 216.87.84.209
```

Create the guest-logon-access policy:

```
ip access-list session guest-logon-access
   user any udp 68 deny
   user any svc-dhcp permit time-range working-hours
   user alias “Public-DNS" svc-dns src-nat pool dynamic-srcnat time-
   range working hours
```

Create the auth-guest-access policy:

```
ip access-list session auth-guest-access
   user any udp 68 deny
   user any svc-dhcp permit time-range working-hours
   user alias “Public-DNS” svc-dns src-nat time-range working-hours
   user any svc-http src-nat pool dynamic-srcnat time-range working-
   hours
   user any svc-https src-nat pool dynamic-srcnat time-range
   working-hours
```

Create the block-internal-access policy:

```
ip access-list session block-internal-access
   user alias “Internal-Network” any deny
```

Create the drop-and-log policy:

```
ip access-list session drop-and-log
   user any any deny log
```

Create the guest-logon role:

```
user-role guest-logon
   session-acl captiveportal position 1
   session-acl guest-logon-access position 2
   session-acl block-internal-access position 3
```

Create the auth-guest role:

```
user-role auth-guest
   session-acl cplogout position 1
   session-acl guest-logon-access position 2
   session-acl block-internal-access position 3
   session-acl auth-guest-access position 4
   session-acl drop-and-log position 5
```

Configure the guest VLAN:

```
vlan 900

interface vlan 900
   ip address 192.168.200.20 255.255.255.0

ip dhcp pool "guestpool"
   default-router 192.168.200.20
   dns-server 64.151.103.120
   lease 0 4 0
   network 192.168.200.0 255.255.255.0
```

Configure captive portal authentication:

```
aaa authentication captive-portal guest-net
   default-role auth-guest
   user-logon
   no guest-logon
```

Modify the guest-logon role:

```
user-role guest-logon
   captive-portal guest-net
```

Configure the AAA profile:

```
aaa profile guest-net
   initial-role guest-logon
```

Configure the guest WLAN:

```
wlan ssid-profile guest-net
   essid guest-net
   opmode opensystem

aaa profile guest-net
   initial-role guest-logon

wlan virtual-ap guest-net
   vlan 900
   ssid-profile guest-net
   aaa-profile guest-net
```

With the appropriate levels of encryption and authentication used, for different users associated and authenticated to the same AP at the same time, the system is completely secured. The unique combination of these security mechanisms and Aruba's Role-Based Access Control (RBAC) gives an Aruba User-Centric Network far more control and granularity of user traffic than simply demanding a particular type of authentication and encryption. This same flexibility gives customers the ability to deploy Remote APs that broadcast Employee SSIDs at a user's home for telecommuting or at another business to conduct a sales demonstration without fear of security breach.

Device Role

Special-purpose device roles are very similar to the guest access role; and most commonly include active RFID tags, voice and video devices. Device roles should be setup to allow them to perform only single functions and to be able to interact only with a known set of IP addresses. For example, a voice device should only be able to run voice protocols such as Session Initiation Protocol (SIP) to the SIP server, Real-Time Transport Protocol (RTP) and basic ICMP commands. Any other uses should result in the device being blacklisted as it is most likely the subject of an impersonation attack.

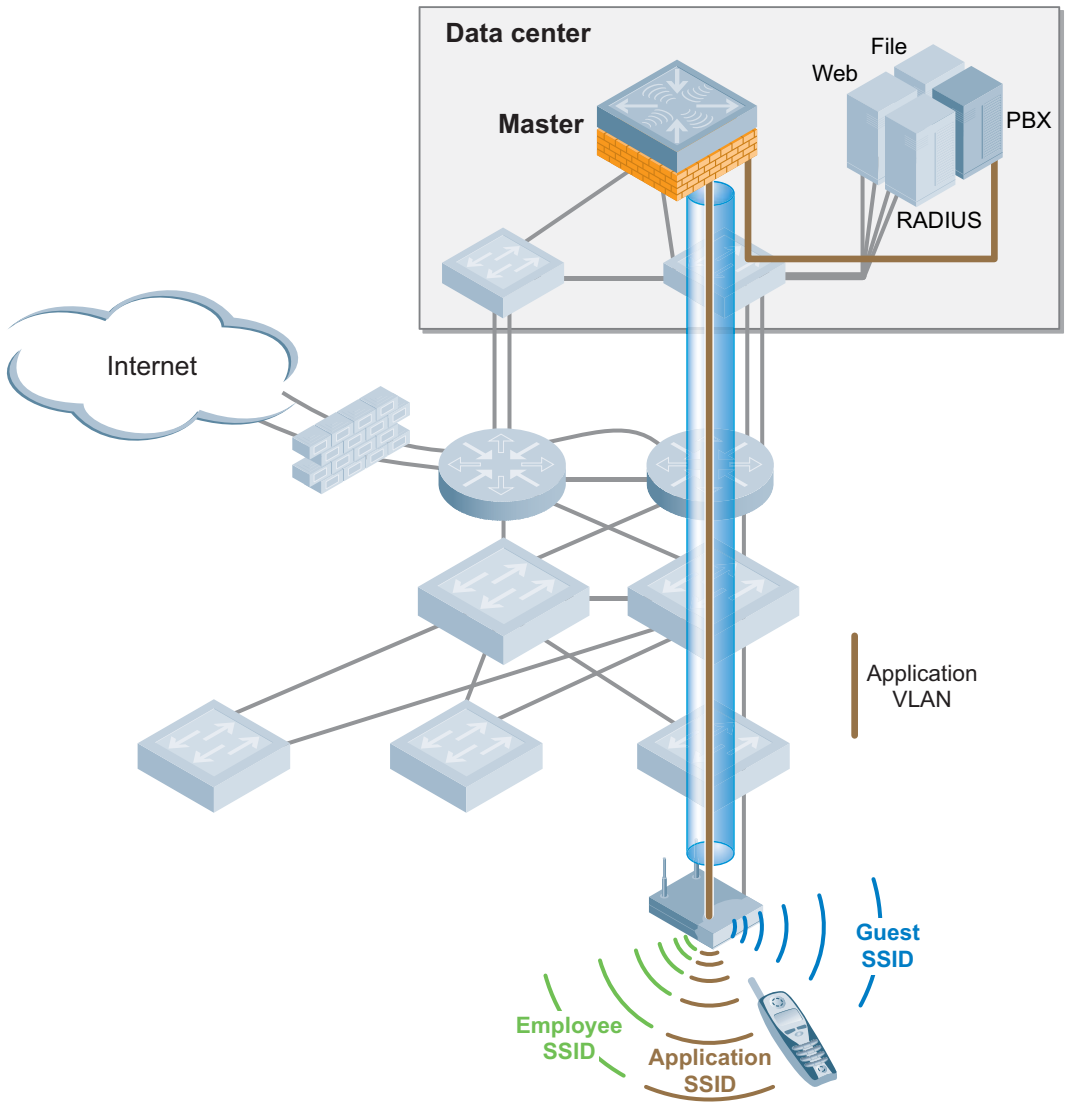

Role Variation by Authentication Method

Role assignment has many options under the umbrella of role derivation. While the system can simply use the role returned from the authentication server, it can also assign a role based on a number of attributes. When a user logs in using WPA2 they receive an open employee role, but when logging in with the same credentials using a captive portal, a reduced role is put into effect. Phones can share the same authentication as a camera but receive a different voice role after completing registration with a SIP server.

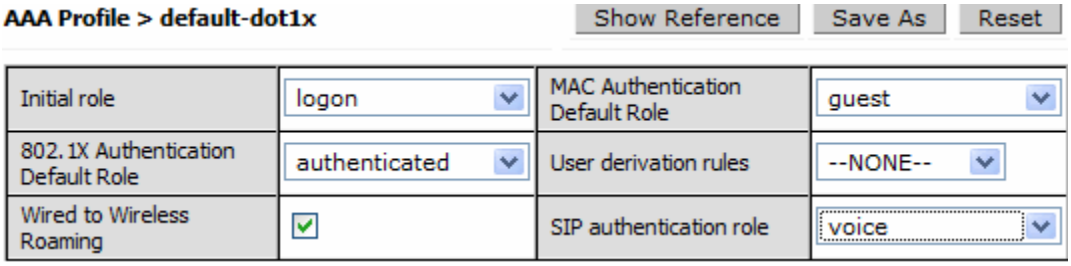

Wireless Intrusion Detection System

Due to the open nature of wireless networks they are prime targets for both unscrupulous individuals and casual hackers that happen to live in the area. To protect against unsanctioned wireless devices, use of Aruba's Wireless Intrusion Detection System (WIDS) software module can automatically detect and defend against wireless attacks and "rogue" APs on the network.

Wireless Attacks

A common wireless network attack is called "man-in-the-middle." During such an attack, a hacker masquerades as a legitimate AP using software on a laptop, and acting as a relay point, fools users and other APs into sending data through the laptop instead of the real AP. The attacker can then eavesdrop on conversations, modify or corrupt data, or run password-cracking routines.

Aruba Access Points monitor the air to detect other wireless stations masquerading as valid APs. When masquerading is detected, appropriate defense mechanisms are put into place. Aruba Mobility Controllers also track unique "signatures" for every wireless client in the network, and if a new station is introduced claiming to be a particular client, but lacks a proper signature, a station impersonation attack is declared.

Advanced Denial of Service (DoS) protection keeps enterprises safe against a variety of other wireless attacks, including association and de-authentication floods, 'honeypots' and AP or station impersonations. Based on location signatures and client classification, Aruba access points will drop illegal requests and generate alerts to notify administrators of the attack. The system will report attacks to network administrators, and take proactive measures to prevent users from falling victim to these attacks.

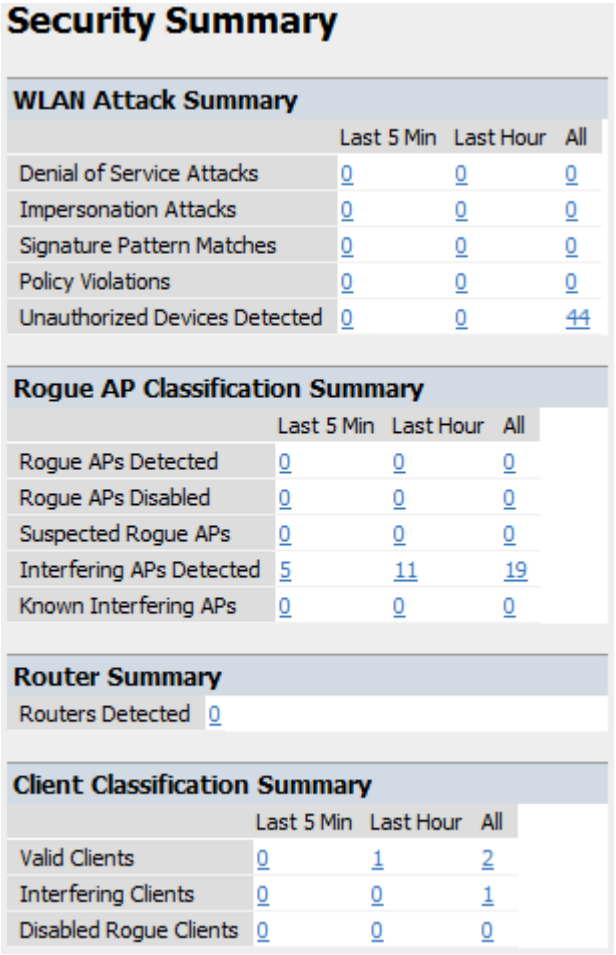

Security Summary

WLAN Attack Summary

	Last 5 Min	Last Hour	All
Denial of Service Attacks	0	0	0
Impersonation Attacks	0	0	0
Signature Pattern Matches	0	0	0
Policy Violations	0	0	0
Unauthorized Devices Detected	0	0	44

Rogue AP Classification Summary

	Last 5 Min	Last Hour	All
Rogue APs Detected	0	0	0
Rogue APs Disabled	0	0	0
Suspected Rogue APs	0	0	0
Interfering APs Detected	5	11	19
Known Interfering APs	0	0	0

Router Summary

Routers Detected	0

Client Classification Summary

	Last 5 Min	Last Hour	All
Valid Clients	0	1	2
Interfering Clients	0	0	1
Disabled Rogue Clients	0	0	0

Rogue APs

There are two types of 'Rogue APs'; one that is not connected to your wired network and one that is. An unconnected Rogue AP could be set up inside your office by a contractor or well-meaning employee to provide wireless service to a small group of users. However, this AP is consuming precious spectrum and potentially creating co-channel interference with authorized enterprise APs in the area. A connected Rogue AP is when an employee or contractor takes a consumer-grade access point and plugs it into a nearby open network port to provide a personal hotspot.

'Rogue Classification' means the process of detecting the presence of a Rogue AP and determining which type it represents. The rogue AP classification algorithms allow the system to accurately differentiate between threatening 'rogue' APs installed on the Local network and nearby 'interfering' APs. Once classified as rogue, these APs can be automatically disabled through both the wireless and wired networks. Administrators are also notified of the presence of rogue devices, along with the precise physical location on a floor plan, so that the rogue device can be removed from the network.

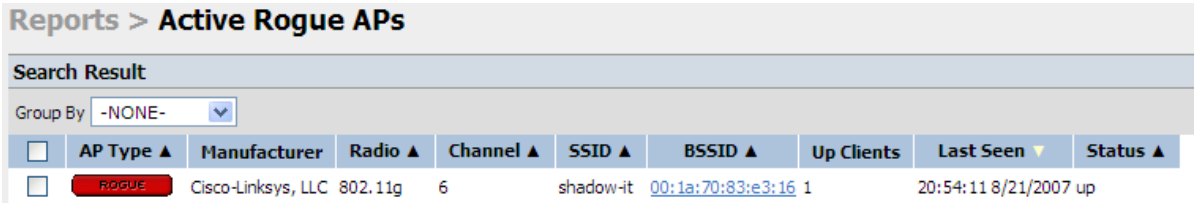

Once classification is complete, 'Rogue Containment' refers to the Aruba system taking active action against the Rogue AP, if it is configured to do so. Aruba wireless systems are often configured to detect but not to automatically contain rogue devices, and can be set to different thresholds. Most organizations do not want to take the risk of having the system inadvertently contain APs that may be legitimately part of another business or home. After detecting and alerting, the network administrator can confirm the rogue APs location, and either mark the AP as "known interfering" or choose to manually contain the AP.

Automatic protection for users is often enabled. This prevents "man-in-the-middle" and other attacks for which signatures are well known. It also detects the use of tools often employed to launch attacks against targets within the enterprise.

Wireless networks break many of the old rules when it comes to designing a network; as a result, new tools are needed to help administrators quickly and effectively deploy and maintain a wireless network. The following sections describe the RF Plan and Adaptive Radio Management (ARM) tools that Aruba provides to help operate the WLAN effectively and efficiently.

RF Plan Tool

RF planning, in the days of autonomous Access Points, was a painful experience at best. It was often a headache to look at a two-dimensional map of AP placements and attempt to determine which channel and power setting should be used. Because early Access Points were extremely expensive and the widest coverage possible from each was needed, a detailed site survey was performed accounting for building wall construction and possible interference sources. As the wireless link becomes the primary connection for major enterprise deployments and the cost of APs has fallen, the need to increase AP density to allow higher throughput has changed site survey process.

Aruba Networks recommends a dense deployment of APs as discussed earlier in this document. This model reduces or eliminates the need for a formal site survey. In general, many professional WLAN designers say it is better to spend less money today on a site survey that provides a limited one-time snapshot of your environment. Instead, it is an industry best practice to use those funds to buy a few more APs that provide improved service and long-term benefit.

The Aruba RF Plan tool easily imports an image of each floor of a building to be covered and suggests AP counts and placement based on the following simple rules:

- Connection needs (speed, coverage, or AP count)
- Redundancy (cell overlap)

The system will suggest a layout that meets the criteria. The user can easily reposition one or more APs to accommodate building features or customize coverage. Once the APs have been placed, the user can examine the predicted RF environment as seen in the image below.

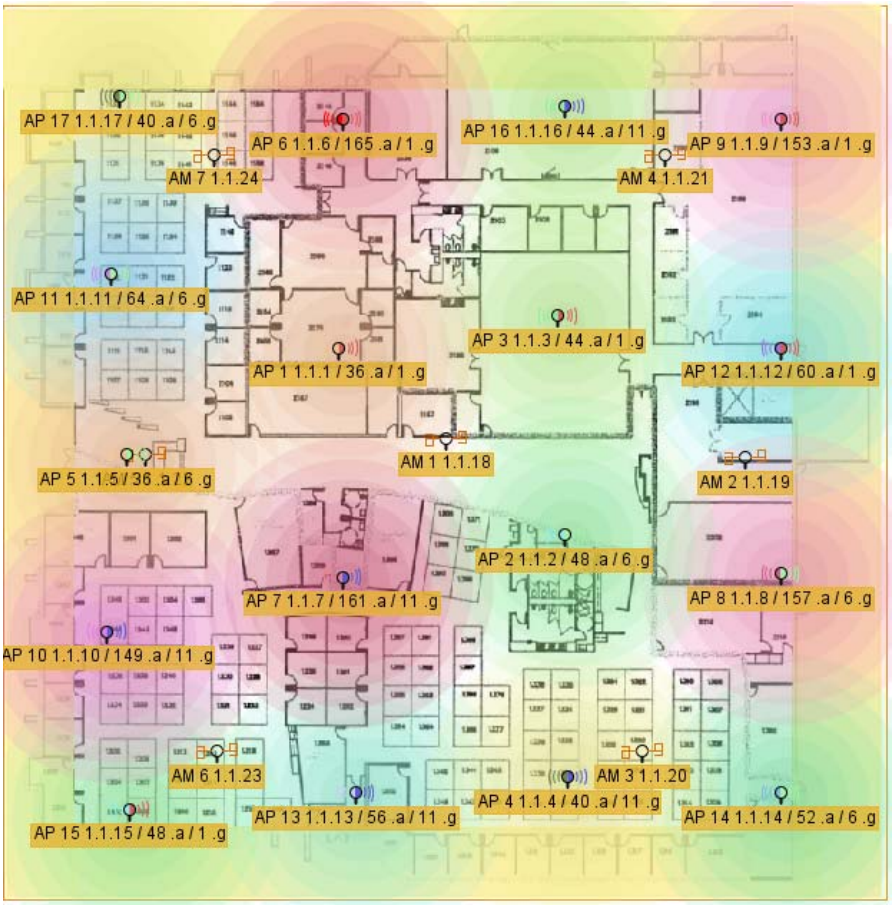

The RF Plan tool is available on the Mobility Controller, Mobility Management System™, and in a standalone version. All of these versions use the same file format, allowing an RF plan to be developed prior to deploying a controller, and then imported into the RF Live application for viewing.

Aruba recommends as a best practice that each customer completes a post-installation site survey to verify that the delivered coverage matches what was predicted. Occasionally, there will be gaps in the coverage or areas that do not meet the predicted signal strength. This is typically due to unexpected construction materials, banks of metal file cabinets or other RF-opaque building features. However, this is infrequent in a dense deployment. If RF gaps are found, a few APs can be adjusted before the cabling team leaves the site, or the spare APs deployed. If there are no coverage gaps, the extra APs can be used as spares for the future or deployed to cover additional areas.

Adaptive Radio Management

Selecting power and channel settings for hundreds or thousands of Access Points across a campus on foot is not something that any administrator would look forward to without software or hardware automation. The RF medium is continuously changing. While today channel 6 may be optimal for a given area, tomorrow the best choice might be channel 1. Something as simple as new construction could alter the RF characteristics of an area, requiring all APs in the local area to be adjusted.

ARM is an RF spectrum management technology that provides a stable, high performing, self healing wireless LAN deployment that does not require Administrator intervention. ARM is a distributed system that involves an AP or AM continuously scanning all of the legal channels within its regulatory domain, and coordinating channel and power settings on all APs using the Mobility Controller. In the Aruba VRD, for Campus Wireless networks, this processing occurs on the Master at the Management layer, and settings are pushed back down to the Locals at the Aggregation layer.

The ARM system handles setting all power and channel setting, including moving the APs to new channel and power settings automatically when appropriate. The network administrator spends no time managing the RF environment even in the case of RF jamming attacks or interference from legitimate wireless sources in the same frequency. The system automatically determines the best settings, and can automatically move away from interference laden channels without any intervention. The ARM system works indoors or outdoors.

Sometimes a channel change is desirable but to do so would interrupt active user sessions. Certain types of devices are more tolerant of such changes than others. Because the Mobility Controller is aware of not only users on the AP but also the type of traffic being sent, the AP can be directed not to change channels for specific client types. It can be set to pause scanning if going off channel will cause unacceptable quality disruption. These features are called Client Aware and Voice Aware Scanning.

Aruba Client Aware Scanning prevents a channel change while clients are associated with a particular AP. This ensures that clients will continue to send and receive data without the AP suddenly switching to a new channel. When all clients have left a particular AP, it is free to change channels. While the client is attached, it will continue to go off channel for scanning at predetermined intervals unless ARM scanning is disabled.

If the customer has chosen not to deploy dedicated Air Monitors, the system can still obtain the information it needs to conduct basic WIPS and ARM functions by 'timesharing' with the data traffic on each AP. This scanning is transparent to data users, but can have detrimental effects on isochronous traffic flows such as voice calls. With Aruba Voice Aware Scanning, in the optional Voice Services Module, the AP will temporarily stop channel scanning when calls are in progress. Because the Mobility Controller is aware of the state of the call, it will pause scanning when the phone goes "off hook" and resume scanning when the call ends.

ARM is typically run in a very aggressive mode when the network is first deployed, allowing the APs to quickly settle their power and channel settings. The following setting should be used to initially settle the network infrastructure, and should be configured in a new profile:

```
Parameter                   Value
---------                   -----
Assignment                  single-band
Client Aware                Disabled
Max Tx Power                30 dBm
Min Tx Power                11 dBm
Multi Band Scan             Disabled
Rogue AP Aware              Disabled
Scan Interval               1 sec
Scanning                    Enabled
Scan Time                   110 msec
Power Save Aware Scan       Disabled
Ideal Coverage Index        5
Acceptable Coverage Index   2
Wait Time                   2 sec
Free Channel Index          25
Backoff Time                120 sec
Error Rate Threshold        50 %
Error Rate Wait Time        30 sec
Noise Threshold             75 -dBm
Noise Wait Time             120 sec
Minimum Scan Time           2 sec
```

Aruba recommends that the above settings should be run for a minimum of one hour, and if possible overnight. Once the network has settled, the following configuration should be used for normal ARM operation. These are the default settings.

```
Parameter                  Value
---------                  -----
Assignment                 single-band
Client Aware               Enabled
Max Tx Power               30 dBm
Min Tx Power               11 dBm
Multi Band Scan            Enabled
Rogue AP Aware             Disabled
Scan Interval              10 sec
Scanning                   Enabled
Scan Time                  110 msec
Power Save Aware Scan      Enabled
Ideal Coverage Index       5
Acceptable Coverage Index  2
Wait Time                  15 sec
Free Channel Index         25
Backoff Time               240 sec
Error Rate Threshold       50 %
Error Rate Wait Time       30 sec
Noise Threshold            75 -dBm
Noise Wait Time            120 sec
Minimum Scan Time          8 sec
```

The wireless administrator should be aware that the aggressive settings may cause connectivity issues for clients when the AP suddenly changes power and channel settings because both client and voice aware features are disabled. However, Aruba recommends going through this phase as a best practice to arrive at an optimal radio configuration.

As more enterprises move from purely data-driven applications and incorporate Voice over IP (VoIP) and streaming video delivery, the demands for quality of service and assured application delivery rise commensurately. Multi-media applications can literally break a wireless LAN not designed for application awareness and automatic flow control.

Aruba's Mobility Controller and Access Point technology encompasses many technological innovations designed specifically to address the needs of toll-quality voice and streaming video applications. The first innovation is application awareness, the ability to detect the type of application and coordinate suitable network conditioning to ensure proper packet delivery. In an Aruba network, this process occurs automatically and requires no user intervention.

Adaptive interference management makes the best use of available 802.11n and 802.11a/b/g radio channels to provide a clear path for signaling. Air-time bandwidth control adjusts available bandwidth on a packet-rate or air-time basis to prevent slow clients from consuming too much channel time.

Mobile clients need to be switched from access point to access point in an efficient and expeditious manner when roaming. Aruba offers instantaneous hand-offs to allow undisturbed roaming without hand-off delays or dropouts, and automatic load balancing to ensure that roaming clients don't slow down access point performance.

The Aruba Mobility Controller comes with basic support for voice communications right out of the box. In addition, Aruba offers an optional Voice Services Module licence which provides advanced features for production voice deployments. The Voice Services Module license is a recommended element of this VRD when voice devices are present.

WMM and QoS

Support for 802.11e and WMM ensures wireless QoS for delay-sensitive applications with mapping between WMM tags and internal hardware queues. Aruba Mobility Controllers also support mapping of 802.1p and IP DiffServ tags to hardware queues for wired-side QoS. Layer-2 QoS capabilities are easily enhanced to Layer-3+ flow management and DiffServ using the add-on Policy Enforcement Firewall module.

The system can automatically take steps such as setting quality of service (QoS) parameters and pausing off channel AP scanning to insure voice and video transmissions receive uninterrupted service. The system can also reprioritize traffic that is set for one service level but actually belongs at a different level.

Quality of Service

Voice communication is sensitive to a number of factors such as end-to-end delay and jitter.

End-to-end delay is the time it takes an analog sound made at the sending device to be reproduced as an analog sound at the receiving device. Sources of delay include codec delay, packetization delay, serialization delay and network delay. For acceptable voice quality, end-to-end delay should be less than 200ms. Aruba Mobility Controllers minimize the network delay component of end-to-end delay – the time for voice packets to cross the network from sending device to receiving device.

Jitter is related to delay, but is the variation in delay between packets. Jitter is typically caused by media contention, buffering, routing changes, and network congestion. Jitter is problematic for constant bit-rate (CBR) traffic like voice because it causes variation in the bit rate, and erodes the voice

call quality. Jitter buffers are used in VoIP networks to smooth out this effect, but they add delay and must be as small as possible. Aruba Mobility Controllers adjust network settings to minimize jitter and maximize voice quality.

Traffic Prioritization

Aruba Mobility Controllers use traffic prioritization as one method to address delay and jitter. Traffic prioritization assures that voice packets have preferential access to the media and are moved ahead of best-effort traffic in buffers during congestion.

Application-based prioritization requires stateful inspection; this capability is a crucial difference between an Aruba solution, and competing wireless solutions. Competing solutions prioritize based on a wireless SSID, meaning that all traffic transmitted on a particular SSID is treated the same. This precludes support for voice applications running on multi-function devices such as laptops or PDAs, since these devices use multiple protocols. Aruba Mobility Controllers contain a policy enforcement firewall that statefully identifies, tracks and dynamically prioritizes traffic based on the application flow, e.g., giving higher priority to a SIP session than an HTTP session, even from the same device.

Network Wide QoS

While the Aruba Mobility Controllers can handle much of the heavy lifting by identifying, properly tagging, and scheduling packets into the network, the rest of the components must also be ready to handle QoS. If the access, distribution, core, and data center switches and routers are only providing best effort delivery, voice quality will suffer. At each level in the network, devices that will be forwarding QoS tagged traffic must be configured to properly prioritize traffic above data and background traffic.

Voice Functionality and Features

Voice Service Module features provide deep visibility into the session, such as, viewing the call progression and voice quality of a SIP based VoIP call. Advanced voice-over-WLAN features such as Call Admission Control (CAC), voice-aware RF management, and voice-specific diagnostics allow the Mobility Controller to deliver enterprise class mobile VoIP capabilities.

Voice-Aware RF Management

As discussed in Chapter 7 on page 55, Aruba's Adaptive Radio Management (ARM) is normally configured to adjust channel and transmit power levels of wireless APs based on nearby interference and other RF conditions. Client devices will react to a channel change by scanning for a new AP and then re-associating as though they were roaming. Most data applications will not be noticeably affected by this action. Voice is highly susceptible to packet loss, however, and a channel change during a voice call will very likely cause packet loss and audible disruption to the call.

Because Aruba Mobility Controllers statefully follow voice protocols, they will not allow a channel change while voice calls are taking place. If a channel change is required, the controller will wait until that AP is no longer handling active voice calls before initiating the channel change.

Call Admission Control

Typical voice codecs (Coder/Decoder) used in VoIP do not consume large amounts of bandwidth. Even with G.711, which uses 64Kbps per call, a typical 802.11b access point could theoretically support nearly fifty simultaneous calls based purely on bandwidth. In practice, the limiting factor is contention for the wireless medium because 802.11 uses a collision-avoidance algorithm that makes timely access to the wireless media a challenge for delay-sensitive devices. Due to this limitation, the number of

simultaneous voice calls handled by a single AP must be limited. This limit varies based on network conditions and handset manufacturer, and is typically provided in a manufacturer's design guidelines.

Call admission control (CAC) is included with the Voice Services Module license. CAC lets the Mobility Controller limit the number of voice calls on an AP, and proactively move voice clients to a less-utilized AP. Aruba Mobility Controllers implement CAC by statefully following voice protocols and being aware of the voice utilization of a given AP. Per-SSID association limits for each AP also prevent a voice device from associating to a dedicated voice SSID when that AP has reached a pre-configured limit.

Comprehensive Voice Management

The Voice Services Module license adds extensive voice management functionality, providing detailed reporting and troubleshooting capabilities. Information is available at a glance via well-organized tables and graphs. Some of the capabilities include:

- Phone number association – SIP devices can be tracked and displayed by their associated phone number.
- Call quality tracking – Automatically calculates, displays and tracks the R-value for each SIP call being processed through the Aruba mobility controller.
- SIP authentication tracking – Tracks the registration of SIP devices with a IP PBX to determine if they are authenticated devices.
- Call detail records (CDRs) – Displays the calls made to or from Wi-Fi clients, including originator, terminator, termination reason, rejected and failed calls, duration, call quality, etc.
- CAC-based real-time information – Quickly determine call density, CAC state, and active calls.

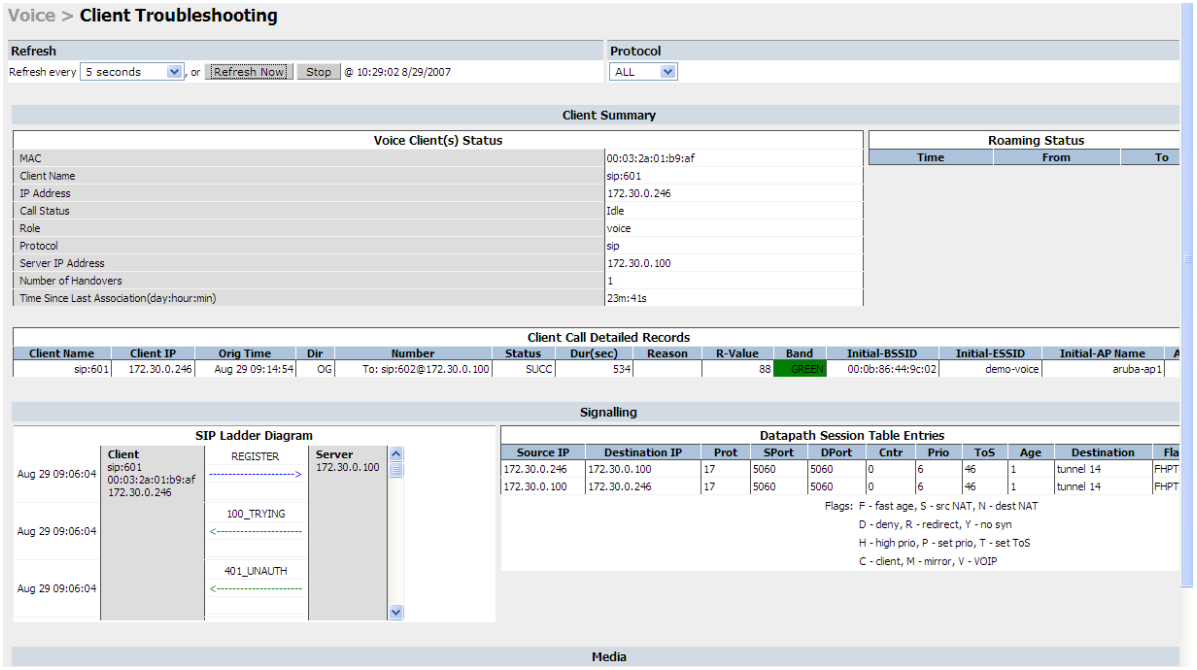

Chapter 9

Controller Clusters and the Mobility Management System™

The Aruba Mobility Management System™ (MMS) is designed to give network administrators the ability to manage the system state and rapidly scope problems for individual users across one or more Master/Local Mobility Controller clusters.

As the network grows beyond a single Master/Local cluster the configuration and troubleshooting of the system grows in complexity. This complexity is increased further if more than a single cluster exists on the same campus as users could easily roam between clusters. To simplify the job of the network administrator, Aruba recommends using the MMS system any time more than one Master/Local Mobility Controller cluster exists in the network.

The MMS product provides a consolidated view of all components and users on the network in a single, flexible console. In addition to the functionality already present in the Mobility Controllers, the MMS adds network-wide configuration, advanced reporting and trending to the Aruba system, allowing network administrators to interface with a single tool for planning, configuration, and troubleshooting.

The Mobility Management System reduces total cost of ownership by automatically discovering and managing hundreds of controllers and thousands of access points and users from a single network operations center.

Centralized configuration management, coupled with the ability to track client devices, identify rouge devices, and plan new deployments and visualize RF coverage patterns with an intuitive, seamless user interface, is a key differentiator. MMS provides a comprehensive suite of applications for planning, configuration, fault and performance management, reporting, RF visualization, and Wi-Fi® device and RFID location tracking for Aruba's User-Centric Networks. This product seamlessly integrates with Aruba's Access Points and Mobility Controllers to support the new paradigm of adaptive wireless LANs, identity-based security, and application continuity.

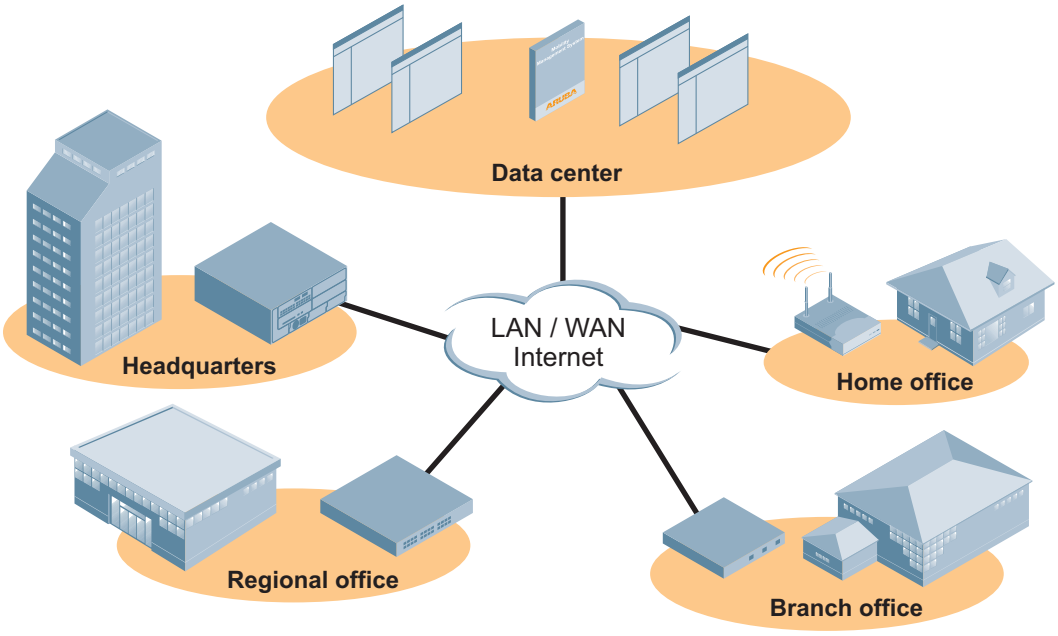

Configuration is handled by the same Profile system discussed n Chapter 6 on page 37. With Mobility Controller clusters grouped on the MMS console, different Master/Local clusters can share the same configuration or have different configurations by cluster. Configuration checkpoints and recovery can be performed, as well as the ability to configure changes but apply them at another time. This flexibility reduces errors by sharing common configuration parameters while preserving the ability to have each cluster running a custom configuration.

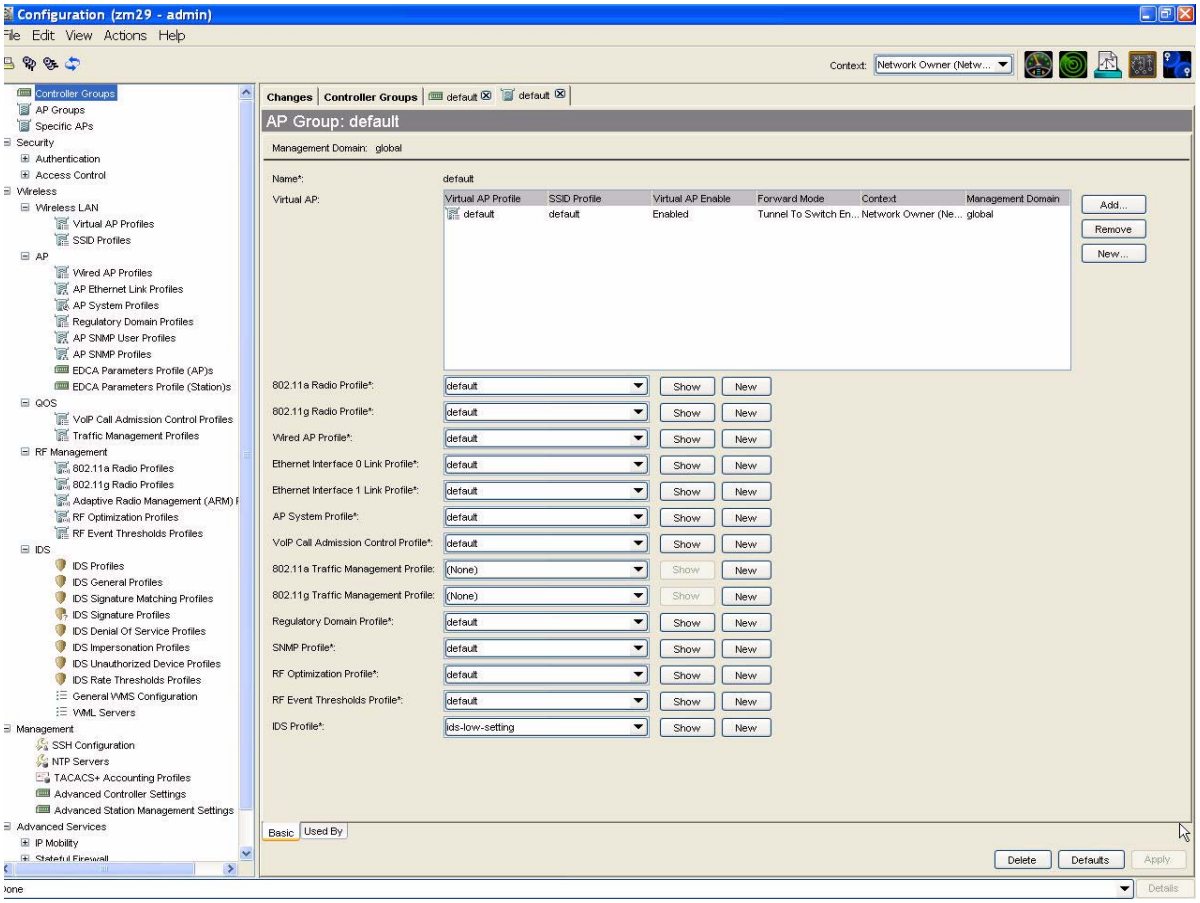

This common configuration capability also eases the administrative burden in creating a Mobility Domain to allow Mobile IP to function across multiple Master/Local clusters. The configuration will be identical and can be pushed to all of the controllers even if they do not share an identical configuration.

The MMS solution will produce a number of standard reports to help with trending and capacity planning, and can be easily configured to do custom reporting. With the built-in hard drive, reports are available for as long as twelve months.

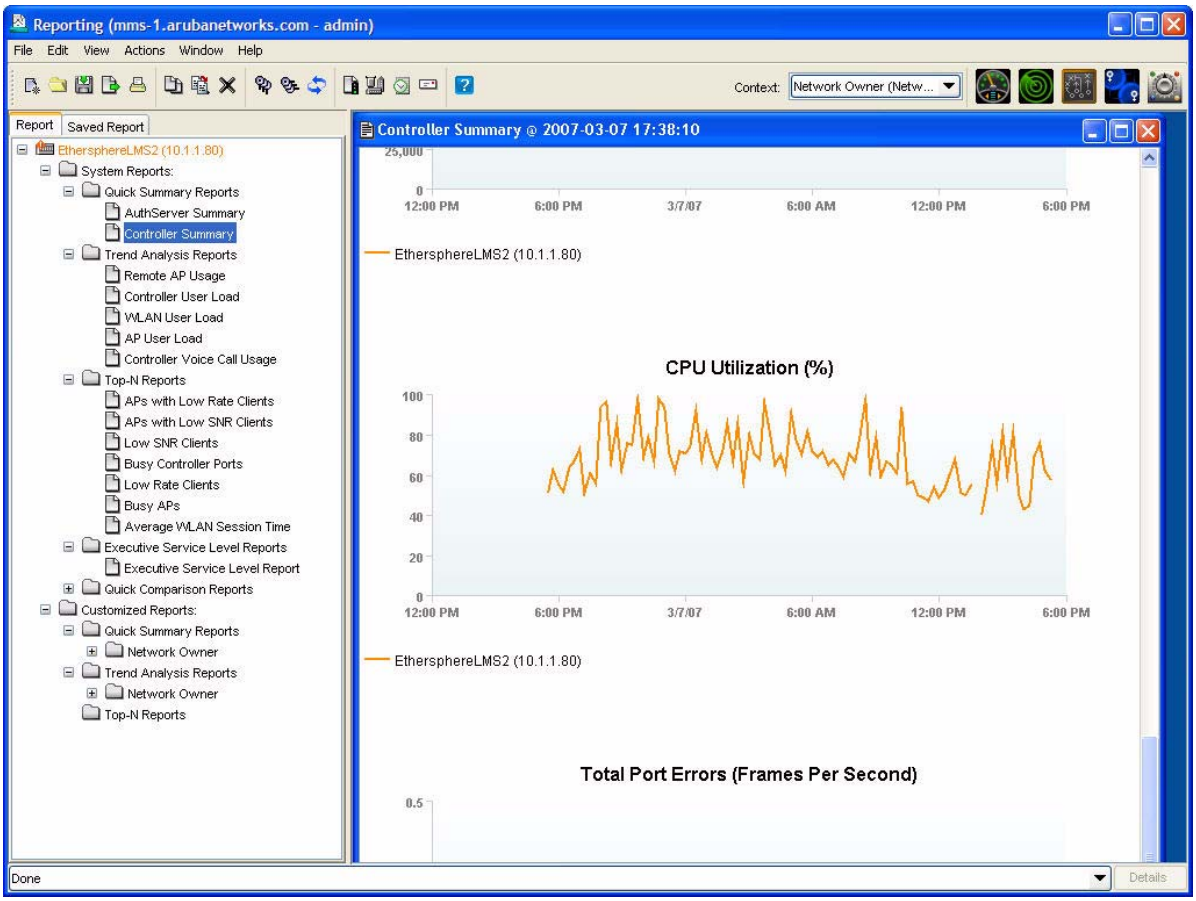

The visualization features of MMS allow the network operations center to quickly view information about the system and its users. The administrator can save searches through the system, allowing them to quickly automate repetitive setup tasks and find the information they need quickly. In addition, the system can be configured to run the searches automatically and email the reports to administrators at configured times.

The same heat maps and location tools available on the controller are also available on the MMS. The location API further extends this capability by allowing 3rd party applications to directly access the system to provide additional custom services. These include RFIDs tags and custom built location applications.

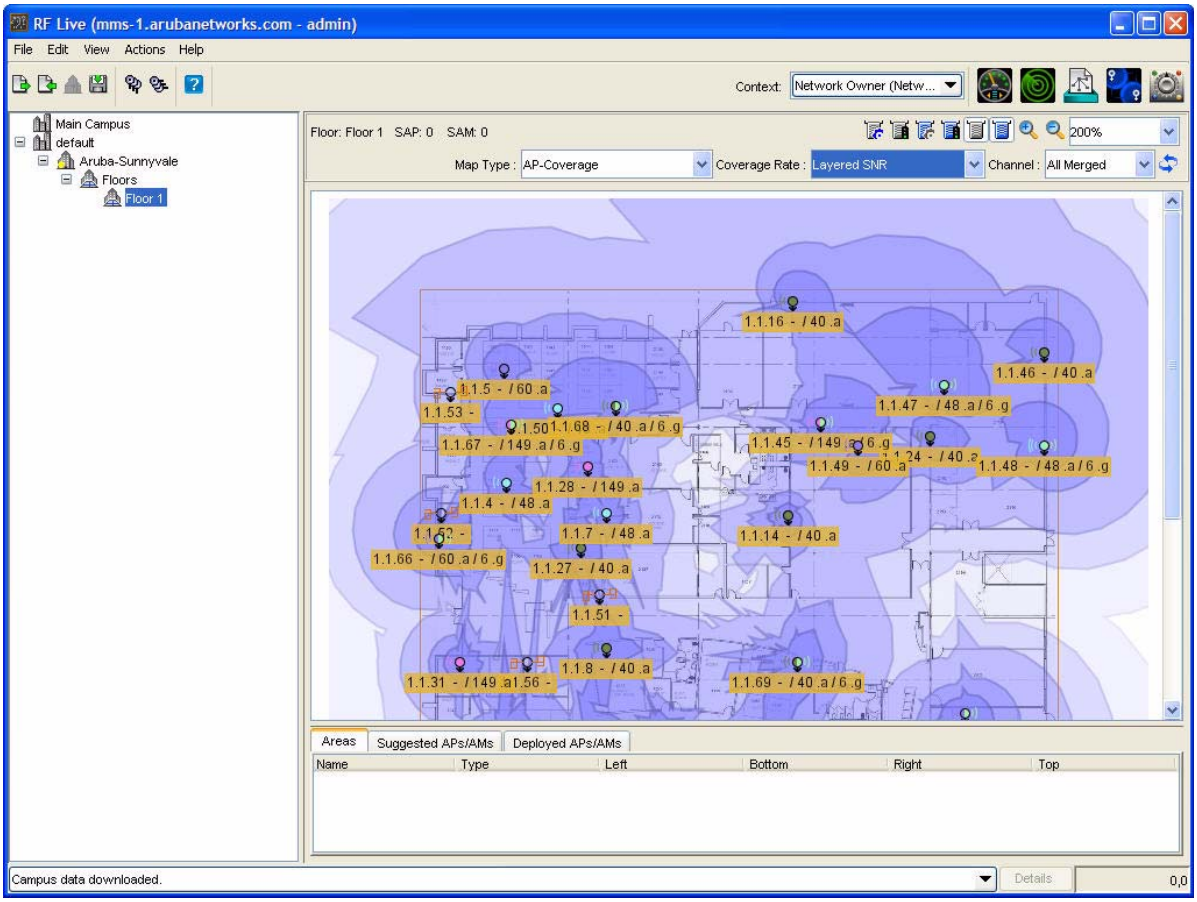

The MMS system uses a Java Web Start user interface; no additional client side software is required. Initial configuration requires setting up SNMPv3 users on all Mobility Controllers and giving the MMS unit the IP information for the Master Controllers in the network. The MMS will then auto-discover the remainder of the network including all Local Controllers and APs/AMs.

Appendix A

Licenses

To extend the base capabilities of ArubaOS, a number of licensed software modules provide additional functionality, including:

Voice Services Module

Delivers standards-based voice over Wi-Fi plus voice control and management innovations enabled by Aruba's application-aware architecture. VSM supports large-scale voice deployments and provides a foundation for fixed mobile convergence (FMC).

Policy Enforcement Firewall

Enforces user-based network access and application priority policies. Policies can be centrally defined and enforced on a per-user basis based on user role and authorization levels. These policies follow users as they move throughout the enterprise network.

Wireless Intrusion Protection

Identifies and protects against malicious attacks on wireless networks, as well as vulnerabilities caused by unauthorized access points and client devices.

Remote Access Point

Extends the enterprise network to small branch offices and home offices having a wired Internet connection. Remote AP software, coupled with any Aruba access point, allows seamless connectivity at home, in a hotel room, or other remote locations.

VPN Server

Extends the mobile enterprise network to large branch offices and individual users over the public Internet, eliminating the need for separate external VPN equipment.

External Services Interface

Delivers a set of control and management interfaces to seamlessly integrate third-party network devices, incremental software modules and services into Aruba's architecture.

xSec

Provides wired and wireless Federal Information Processing Standard (FIPS) 140-2 validated encryption technology designed for high-security government networks.

Appendix B

WLAN Extension with Remote AP

Remote Access Point (RAP) solutions involve configuring a standard thin access point to provide a customer-defined level of service to the user by tunneling securely back to the corporate network over a wide area network. The WAN may be either be a private network such as a frame relay or MPLS network, or a public network such as a residential or commercial broadband Internet service. The same SSIDs, encryption, and authentication that exist on the corporate network are present on the RAP, or the administrator can choose to enable just a subset of the functionality of campus-connected APs. The Remote AP is a licensed feature, with each Remote AP requiring a separate license.

For telecommuter or home-office applications, an Aruba RAP is much more than a simple home wireless device. It is instead an extension of all of services available on the corporate network including voice and video in a similar fashion to a branch office but with fewer configuration headaches. For instance, the user's laptop will automatically associate with the RAP just as it would in the corporate network, and allows for centralized management of a truly mobile edge. Dual-mode voice devices can place and receive calls.

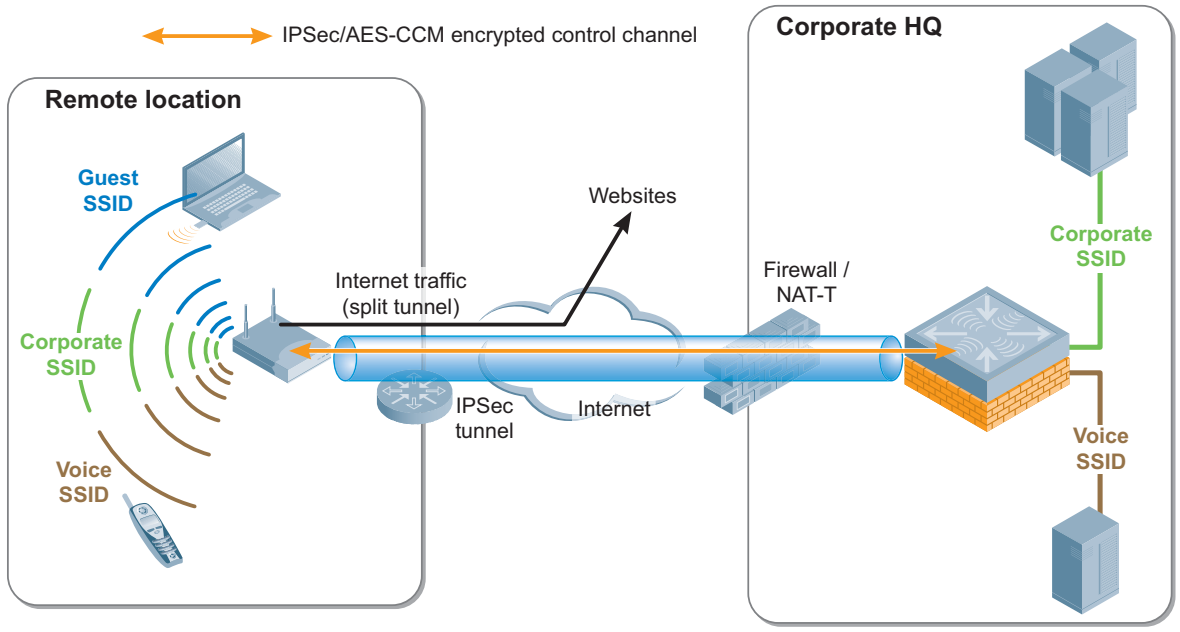

The feature integration of the RAP functions into both the Mobility Controller and thin AP as an end-to-end system is critical to having a solution that is both technologically and cost effective. By integrating authentication, encryption, firewall, and QoS features the network administrator has a single point of troubleshooting and maintenance. This reduces both initial capital expenditure as well as ongoing maintenance costs.

A much larger benefit that comes with this solution is transparent security. The RAP provides a solution that does not add any additional burden to the user beyond their regular login credentials. They simply see connectivity to the home office the same as it is when they are in the office. There is nothing new to remember to do, no tokens to lose, and no mistakes in connecting.

To connect to the Mobility Controller that is inside the corporate network, the Remote AP uses NAT Transversal (NAT-T) to connect through the corporate firewall to the Mobility Controller.

The AP itself should be configured to perform split tunneling. In this configuration the AP will perform decryption of wireless traffic and bridge traffic locally when it is bound for a non-corporate address, and re-encrypt the session using IPSec from the RAP to the corporate controller. The connection to the Internet is protected with the same stateful firewall available on the Mobility Controllers to protect the user from inbound traffic.

Appendix C

Alternative Deployment Architectures

This Campus Wireless LAN Reference Architecture represents a large scale, highly available WLAN deployment model in a single large campus environment. While this is the recommended deployment for this environment, there are other reference architectures that are considered best practices at different scales, and for different types of customers. Aruba has identified four specific reference architecture models in addition to the Campus WLAN that are commonly deployed by our customers.

- Small Deployment (No Redundancy)
- Medium Deployment (1:1 Redundancy)
- Branch Office (N+1 Redundancy)
- Pure Remote Access (1:1 Redundancy)

Each of these scenarios will be covered briefly in the following sections. All of these architectures include a concept of an Aggregation layer and a Management layer as well as discussion of available redundancy options and controller placement. The recommendations for VLANs, profiles, and AP placements are the same as for the Campus WLAN for the most part.

Small Network Deployment

In a small office the network will look much like the Proof-of-Concept design in Chapter 3 on page 15, with a single Mobility Controller and a limited number of APs and AMs. This type of WLAN deployment is typically specified where the WLAN is a convenience network that is not relied upon as the primary connection by users and voice services are not present.

In this scenario the Management layer and Aggregation layer are contained within the same controller, and there is no redundancy. Should the Mobility Controller become unreachable, all APs will go down and the wireless network will be unavailable until the Mobility Controller is once again online.

In this scenario, the Mobility Controller is typically deployed in either the network data center or in the wiring closet. The choice is typically dependent on the physical size of the network and Power-over-Ethernet (PoE) requirements. In a larger physical network that is deploying WLAN in hotspots, the Mobility Controller should be located in the data center. In very small networks where PoE from the controller will also power the APs, the Mobility Controller should be located in the wiring closet. Both options are shown in the following diagrams.

Figure 1 *Mobility Controller located in the network data center*

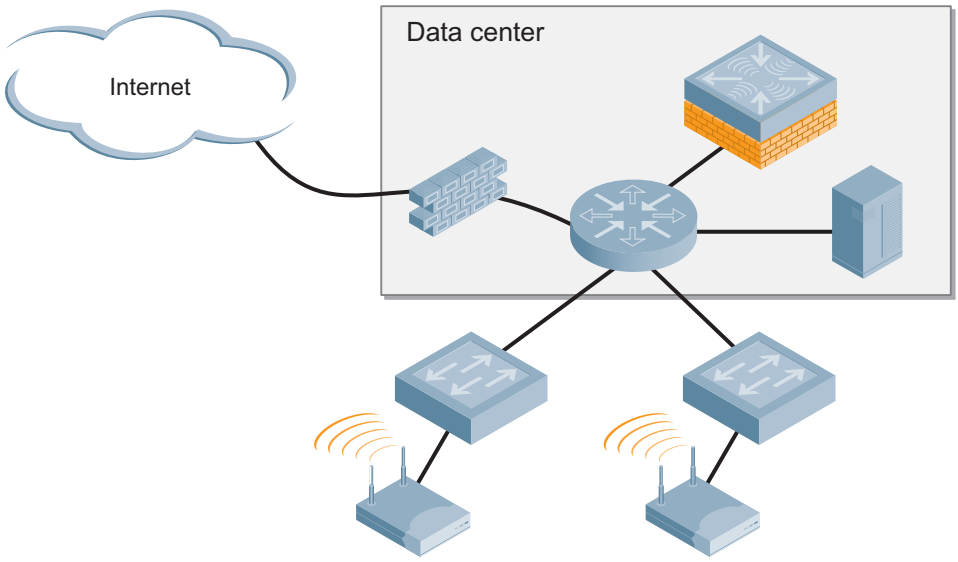

Figure 2 *Mobility Controller located in the common wiring closet (IDF)*

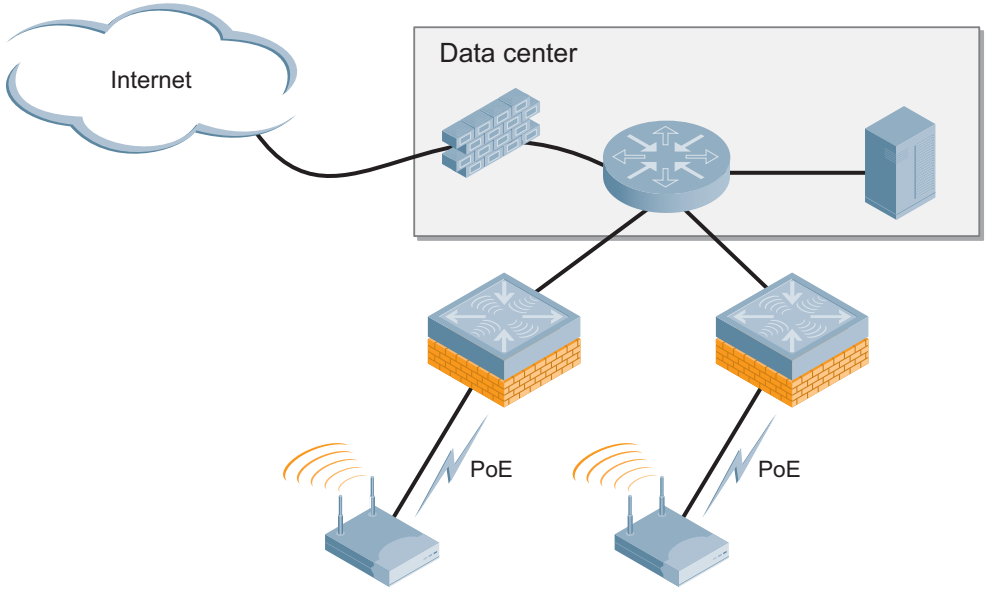

The controllers of choice in this deployment are dependent on AP count and PoE requirements. For small offices requiring PoE, the MC-800 or MC-2400 are both capable of providing power for APs. In offices where PoE is not required, the MMC-3200, MMC-3400, and MMC-3600 series controllers provide a range of AP scaling without the additional costs associated with PoE.

Medium Network Deployment

A medium size network is different from a small network in that the network has moved into general production use and controller redundancy is required. At this point, PoE is no longer provided by the Mobility Controller, and a reference architecture for this deployment model would rely on access layer switches for this function. Additionally, two Mobility Controllers exist in the network to provide high availability.

Redundancy in this model is handled via Master redundancy, with both controllers acting as a Master Mobility Controller. One controller will be in standby, and should be deployed such that it is not serviced by the same power and data connections as the primary Master. Both Mobility controllers are typically deployed in the same data center. As with the Small Network Deployment, the Management and Aggregation layer are coresident in the same production controllers.

Figure 3 *Redundant Master Mobility Controllers deployed in the network data center*

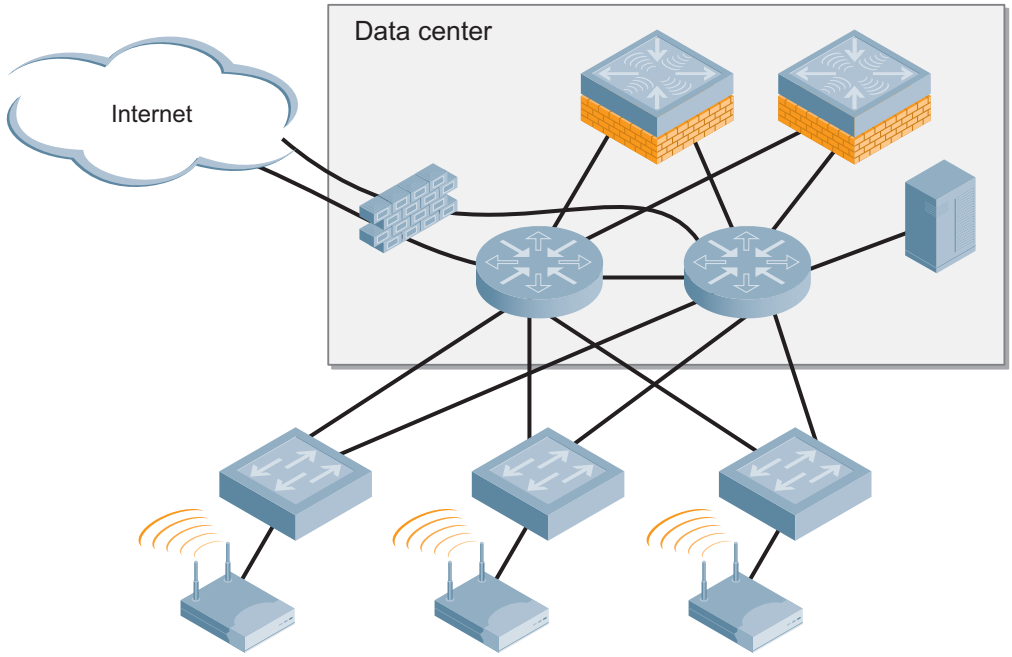

The typical controllers that would be selected for this type of deployment are the MMC-6000 series controllers or the Multiservice Module embedded in the MMC-6000 chassis based controller. The chassis approach offers the advantage of redundant power supplies for greater reliability. The choice should be made based on the size of the network and the expected growth patterns.

Branch Office Deployment

Many large organizations have remote sites that do not have a local IT staff. It is common that such locations have high bandwidth, high availability links that exist to a central data center. For these deployments, wireless redundancy is typically handled across the WAN link to the central data center instead of placing a redundant controller onsite. There may be some type of on-demand backup connectivity in the event of a primary WAN link failure.

The recommended reference architecture for this deployment model would include a single Local Controller that is deployed at each site, with the Master Controller in the central data center acting as the redundant controller. This redundancy model is termed “N+1” because the central controller is intended to provide continuity for the failure of just a few remote controllers at any given time. It requires that APs do not receive power via PoE from the Local Controller and that the Local Controller is not the default gateway for the local site.

The Master Controller is the backup for all Local controllers, and it should be scaled such that a number of sites could potentially encounter issues and remain operational with APs terminating on the Master Controller. The Master must be licensed according to the maximum number of APs and users expected to fail over at any one time. The Master Controller should be deployed in a redundant pair at the central data center to ensure availability. The Management and Aggregation layers are coresident in the data center controllers.

Figure 4 *A single Master Mobility Controller pair backs up all Local Mobility Controllers*

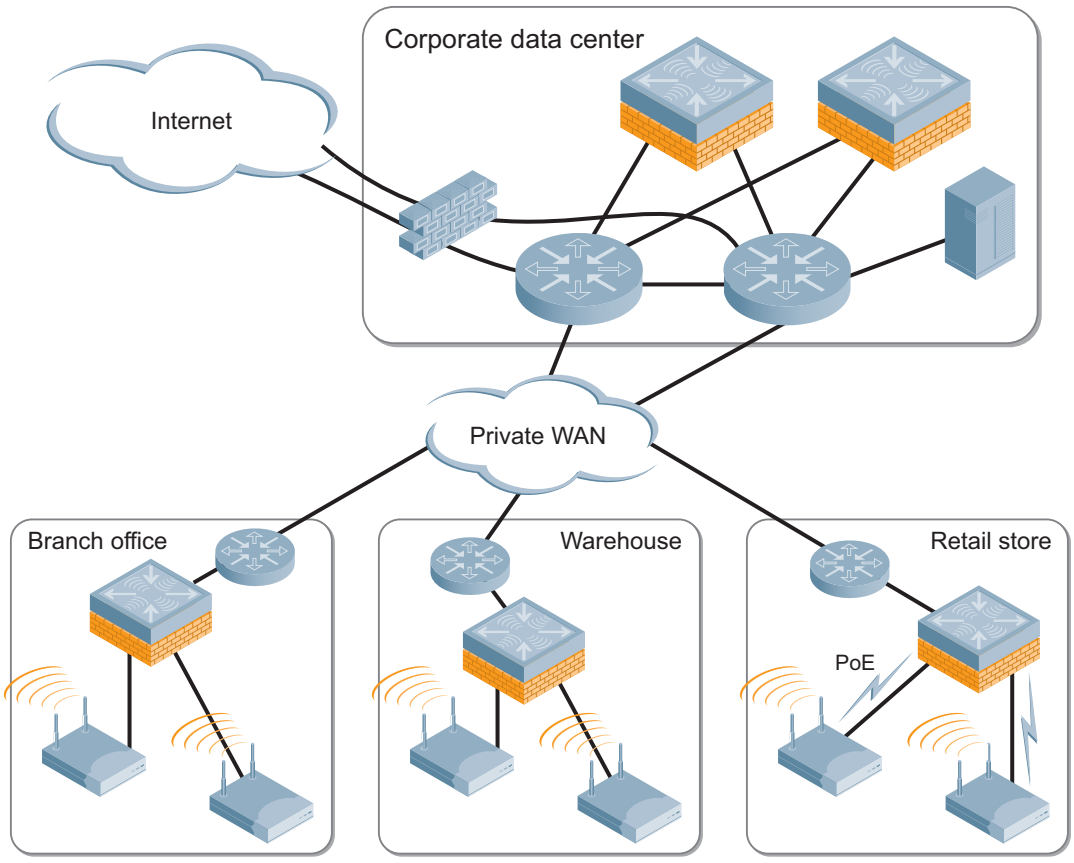

In this scenario the Local Controller a customer would select will typically be a MMC-3000 series controller. The Master Controllers should be MMC-6000 chassis systems to provide the greatest number of AP and users available on the backup system. The chassis should be deployed such that there are no common power or data connections for the Mobility Controllers.

Pure Remote Access Deployment

In some instances, the scale of the Remote AP solution or security requirements dictate that the internal Mobility Controllers serving campus users should not be used for termination of wide-area APs. Typically this means that dedicated Mobility Controllers are placed in the Demilitarized Zone (DMZ) of the network. These Mobility Controllers are solely responsible for terminating RAP and IPSec connections from users.

In this scenario it is important that controllers be highly available because Remote AP functionality is delivered as an "always-on" service. The controllers in this reference architecture are often deployed in Master/Local clusters of two controllers using Active-Active redundancy. These devices also typically straddle the corporate firewall to provide access back into the enterprise just as a typical IPSec concentrator would.

Figure 5 *Remote access Mobility Controllers sit in the network DMZ*

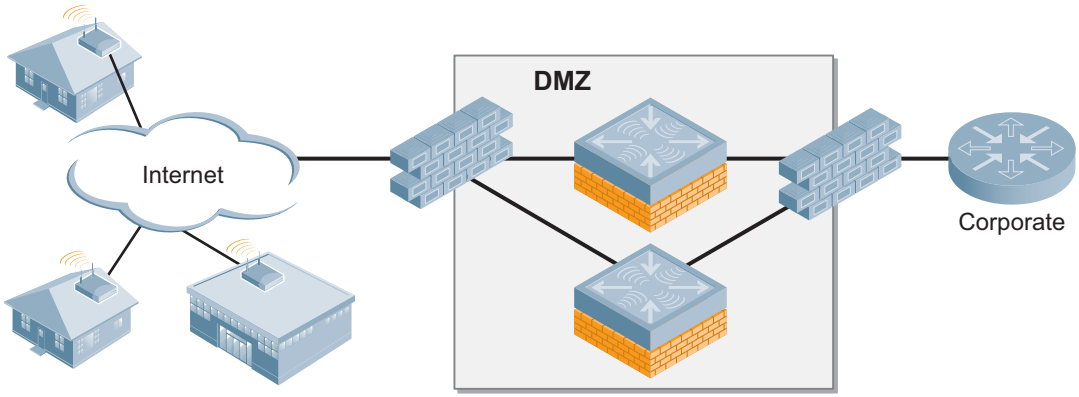

When using stand alone remote access Mobility Controllers it is highly advised that MMS be used in the network to provide configuration. This ensures that all controllers receive the same user roles and firewall policy. This is critical to ensure that the user experiences the same privilege level on the Remote AP as they would on the corporate WLAN.